# Quantum Information and Computation

**Jeffrey Bub**

Department of Philosophy, University of Maryland, College Park, MD 20742*

February 1, 2008

**Abstract**

This Chapter deals with theoretical developments in the subject of quantum information and quantum computation, and includes an overview of classical information and some relevant quantum mechanics. The discussion covers topics in quantum communication, quantum cryptography, and quantum computation, and concludes by considering whether a perspective in terms of quantum information sheds new light on the conceptual problems of quantum mechanics.

*Keywords:* quantum information; quantum computation; quantum teleportation; quantum cryptography; entanglement; quantum measurement; quantum foundations

## Contents



*E-mail address:* jbub@umd.edu

# 1 Introduction

The subject of quantum information has its roots in the debate about conceptual issues in the foundations of quantum mechanics.

The story really begins with the dispute between Einstein and Bohr about the interpretation of quantum states, in particular the interpretation of so-called 'entangled states,' which exhibit peculiar nonlocal statistical correlations for widely separated quantum systems. See, for example, [Bohr, 1949, p. 283] and Einstein's reply in the same volume [Schilpp, 1949]. Einstein took the position that quantum mechanics is simply an incomplete theory. On the basis of a certain restricted set of correlations for a pair of systems in a particular entangled state, Einstein, Podolsky, and Rosen (EPR) argued in a seminal paper [Einstein *et al.*, 1935] that the phenomenon of entanglement conflicts with certain basic realist principles of separability and locality that all physical theories should respect, unless we regard quantum states as incomplete descriptions.

Bohr's view, which he termed 'complementarity,' eventually became entrenched as the orthodox Copenhagen interpretation, a patchwork of reformulations by Heisenberg, Pauli, von Neumann, Dirac, Wheeler, and others. (For a discussion, see [Howard, 2004] and Landsman, this vol., ch. 5.) As Pauli put it in correspondence with Max Born

[Born, 1971, p. 218], a ‘detached observer’ description of the sort provided by classical physics is precluded by the nature of quantum phenomena, and a quantum description of events is as complete as it can be (in principle). Any application of quantum theory requires a ‘cut’ between the observer and the observed, or the macroscopic measuring instrument and the measured system, so that the description is in a certain sense contextual, where the relevant context is defined by the whole macroscopic experimental arrangement. So, for example, a ‘position measurement context’ provides information about position but excludes, in principle, the possibility of simultaneously obtaining momentum information, because there is no fact of the matter about momentum in this context: the momentum value is indeterminate. The Copenhagen interpretation conflicts with Einstein’s realism, his ‘philosophical prejudice,’ as Pauli characterized it in a letter to Born [Born, 1971, p. 221], that lies at the heart of the dispute between Einstein and Bohr about the significance of the transition from classical to quantum mechanics.

The 1990’s saw the development of a quantum theory of information, based on the realization that entanglement, rather than being a minor source of embarrassment for physics that need only concern philosophers, can actually be exploited as a nonclassical communication channel to perform information-processing tasks that would be impossible in a classical world. In a two-part commentary on the EPR paper, Schrödinger [1935, p. 555] identified entanglement as ‘*the* characteristic trait of quantum theory, the one that enforces its entire departure from classical lines of thought.’ This has led to an explosive surge of research among physicists and computer scientists on the application of information-theoretic ideas to quantum computation (which exploits entanglement in the design of a quantum computer, so as to enable the efficient performance of certain computational tasks), to quantum communication (new forms of ‘entanglement-assisted’ communication, such as quantum teleportation), and to quantum cryptography (the identification of cryptographic protocols that are guaranteed to be unconditionally secure against eavesdropping or cheating, by the laws of quantum mechanics, even if all parties have access to quantum computers).

Some milestones: Bell’s analysis [1964] turned the EPR argument on its head by showing that Einstein’s assumptions of separability and locality, applicable in classical physics and underlying the EPR incompleteness argument, are incompatible with certain quantum statistical correlations (not explicitly considered by EPR) of separated systems in EPR-type entangled states. Later experiments [Aspect *et al.*, 1981, Aspect *et al.*, 1982] confirmed these nonclassical correlations in set-ups that excluded the possibility of any sort of physically plausible, non-superluminal, classical communication between the separated systems.

In the 1980s, various authors, e.g., Wiesner, Bennett, and Brassard [Wiesner, 1983, Bennett and Brassard, 1984, Bennett *et al.*, 1982] pointed out that one could exploit features of the measurement process in quantum mechanics to thwart the possibility of undetected eavesdropping in certain cryptographic procedures, specifically in key distribution—a procedure where two parties, Alice and Bob, who initially share no information end up each holding a secret random key which can be used to send encrypted messages between them. No third party, Eve, can obtain any information about the communications between Alice and Bob that led to the establishment of the key, without Alice and Bob becoming aware of Eve’s interference, because Eve’s measurements necessarily disturb the quantum states of the systems in the communication channel.

Bennett [1973] showed how to make a universal Turing machine reversible for any computation, a required step in the design of a quantum computer that evolves via unitary (and hence reversible) state transformations, and Benioff [1980] developed Hamiltonian models for computer computers. Feynman [1982] considered the problem of efficiently simulating the evolution of physical systems using quantum resources (noting that the classical simulation of a quantum process would be exponentially costly), which involves the idea of a quantum computation, but it was Deutsch [1985, 1989] who characterized the essential features of a universal quantum computer and formulated the first genuinely quantum algorithm.

Following Duetsch's work on quantum logic gates and quantum networks, several quantum algorithms were proposed for performing computational tasks more efficiently than any known classical algorithm, or in some cases more efficiently than any classical algorithm. The most spectacular of these is Shor's algorithm [1994, 1997] for finding the two prime factors of a positive integer $N = pq$, which is exponentially faster than the best-known classical algorithm. Since prime factorization is the basis of the most widely used public key encryption scheme (currently universally applied in communications between banks and commercial transactions over the internet), Shor's result has enormous practical significance.

In the following, I present an account of some of the theoretical developments in quantum information, quantum communication, quantum cryptography, and quantum computation. I conclude by considering whether a perspective in terms of quantum information suggests a new way of resolving the foundational problems of quantum mechanics that were the focus of the debate between Einstein and Bohr.

My discussion is heavily indebted to Michael Nielsen and Isaac Chuang's illuminating and comprehensive *Quantum Computation and Quantum Information* [2000], and to several insightful review articles: 'The Joy of Entanglement' by Sandu Popescu and Daniel Rohrlich [1998], 'Quantum Information and its Properties' by Richard Jozsa [1998], and 'Quantum Computing' by Andrew Steane [1998].

# 2 Classical Information

## 2.1 Classical Information Compression and Shannon Entropy

In this section, I review the basic elements of classical information theory. In §2.1, I introduce the notion of the Shannon entropy of an information source and the fundamental idea of information compression in Shannon's source coding theorem (or noiseless channel coding theorem). In §2.2, I define some information-theoretic concepts relevant to Shannon's noisy channel coding theorem.

The classical theory of information was initially developed to deal with certain questions in the communication of electrical signals. Shannon's ground-breaking paper 'A Mathematical Theory of Communication' [Shannon, 1948] followed earlier work by people like Nyquist [1924] and Hartley [1928] in the 1920s. The basic problem was the representation of messages, selected from an ensemble generated by a stochastic process at the message source, in such a way as to ensure their efficient transmission over an electrical circuit such as a noisy telegraph wire.

A communication set-up involves a transmitter or source of information, a (possibly noisy) channel, and a receiver. The source produces messages in the form of sequences of symbols from some alphabet, which Shannon represented mathematically as sequences of values of independent, identically distributed random variables. In later idealizations, the source is represented as *stationary*, in the sense (roughly) that the probability of any symbol (or $n$-tuple of symbols) appearing at any given position in a (very long) sequence, when that position is considered with respect to an ensemble of possible sequences, is the same for all positions in the sequence, and *ergodic*, in the sense that this 'ensemble average' probability is equal to the 'time average' probability, where the time average refers to the probability of a symbol (or $n$-tuple of symbols) in a given (very long) sequence.

The fundamental question considered by Shannon was how to quantify the minimal physical resources required to store messages produced by a source, so that they could be communicated via a channel without loss and reconstructed by a receiver. Shannon's source coding theorem (or noiseless channel coding theorem) answers this question.

To see the idea behind the theorem, consider a source that produces long sequences (messages) composed of symbols from a finite alphabet $a_1, a_2, \ldots, a_k$, where the individual symbols are produced with probabilities $p_1, p_2, \ldots, p_k$. A given sequence of symbols is represented as a sequence of values of independent, identically distributed, discrete random variables $X_1, X_2, \ldots$. A *typical* sequence of length $n$, for large $n$, will contain close to $p_i n$ symbols $a_i$, for $i = 1, \ldots, n$. So the probability of a sufficiently long typical sequence (assuming independence) will be:

$$p(x_1, x_2, \ldots, x_n) = p(x_1)p(x_2)\ldots p(x_n) \approx p_1^{p_1 n} p_2^{p_2 n} \ldots p_k^{p_k n}. \tag{1}$$

Taking the logarithm of both sides (conventionally, in information theory, to the base 2) yields:

$$\log p(x_1, \ldots, x_n) \approx n \sum_i p_i \log p_i := -nH(X) \tag{2}$$

where $H(X) := -\sum_i p_i \log p_i$ is the *Shannon entropy* of the source.

We can think about information in Shannon's sense in various ways. We can take $-\log p_i$, a decreasing function of $p_i$ with a minimum value of 0 when $p_i = 1$ for some $i$, as a measure of the information associated with identifying the symbol $a_i$ produced by an information source. Then $H(X) = -\sum_i p_i \log p_i$ is the average information gain, or the expectation value of the information gain associated with ascertaining the value of the random variable $X$. Alternatively, we can think of the entropy as a measure of the amount of uncertainty about $X$ before we ascertain its value. A source that produces one of two distinguishable symbols with equal probability, such as the toss of a fair coin, is said to have a Shannon entropy of 1 bit: ascertaining which symbol is produced, or reducing one's uncertainty about which symbol is produced, is associated with an amount of information equal to 1 bit.[1] If we already know which symbol will be produced (so the probabilities are 1 and 0), the entropy is 0: there is no uncertainty, and no information gain.

[1] Note that the term 'bit' (for 'binary digit') is used to refer to the basic unit of classical information in terms of Shannon entropy, and to an elementary two-state classical system considered as representing the possible outputs of an elementary classical information source.

Since

$$p(x_1, \ldots, x_n) = 2^{-nH(X)} \tag{3}$$

for sufficiently long typical sequences, and the probability of all the typical $n$-length sequences is less than 1, it follows that there are at most $2^{nH(X)}$ typical sequences. In fact, if the $p_i$ are not all equal, the typical sequences comprise an exponentially small set $T$ (of equiprobable typical sequences) in the set of all sequences as $n \to \infty$, but since the probability that the source produces an atypical sequence tends to zero as $n \to \infty$, the set of typical sequences has probability close to 1. So each typical $n$-sequence could be encoded as a distinct binary number of $nH(X)$ binary digits or bits before being sent through the channel to the receiver, where the original sequence could then be reconstructed by inverting the 1–1 encoding map. (The reconstruction would fail, with low probability, only for the rare atypical sequences, each of which could be encoded as, say, a string of 0's.)

Notice that if the probabilities $p_i$ are all equal ($p_i = 1/k$ for all $i$), then $H(X) = \log k$, and if some $p_j = 1$ (and so $p_i = 0$ for $i \neq j$), then $H(X) = 0$ (taking $0 \log 0 = \lim_{x \to 0} x \log x = 0$). It can easily be shown that:

$$0 \leq H(X) \leq \log k. \tag{4}$$

If we encoded each of the $k$ distinct symbols as a distinct binary number, i.e., as a distinct string of 0's and 1's, we would need binary numbers composed of $\log k$ bits to represent each symbol ($2^{\log k} = k$). So Shannon's analysis shows that messages produced by a stochastic source can be compressed, in the sense that (as $n \to \infty$ and the probability of an atypical $n$-length sequence tends to zero) $n$-length sequences can be encoded without loss of information using $nH(X)$ bits rather than the $n \log k$ bits required if we encoded each of the $k$ symbols $a_i$ as a distinct string of 0's and 1's: this *is* a compression, since $nH(X) < n \log k$ except for equiprobable distributions.

More precisely, let $\overline{X} = \frac{1}{n}(X_1 + X_2 + \ldots + X_n)$, where $X_1, X_2, \ldots, X_n$ are $n$ independent and identically distributed random variables with mean $< X >$ and finite variance. The weak law of large numbers tells us that, for any $\epsilon, \delta > 0$,

$$\Pr(|\overline{X} - < X > | \geq \delta) < \epsilon \tag{5}$$

for sufficiently large $n$.

Now consider a random variable $X$ that takes values $x$ in an alphabet $\mathcal{X}$ with probabilities $p(x) = \Pr(X = x), x \in \mathcal{X}$. [2] Let

$$Z = -\log p(X) \tag{6}$$

be a function of $X$ that takes the value $-\log p(x)$ when $X$ takes the value $x$. Then

$$< Z > = -\sum_{x \in \mathcal{X}} p(x) \log p(x) = H(X) \tag{7}$$

[2]Note that $p(x)$ is an abbreviation for $p_X(x)$, so $p(x)$ and $p(y)$ refer to two different random variables. The expression $\Pr(X \in S) = \sum_{x \in S} p(x)$ denotes the probability that the random variable $X$ takes a value in the set $S$, and $\Pr(X = x)$ denotes the probability that $X$ takes the value $x$. The expression $p(x_1, x_2, \ldots, x_n)$ denotes the probability that the sequence of random variables $X_1, X_2, \ldots, X_n$ takes the sequence of values $(x_1, x_2, \ldots, x_n)$. The discussion here follows Cover and Thomas [1991] and I use their notation.

and for a sequence of $n$ independent and identically distributed random variables $X_1, X_2, \ldots, X_n$:

$$\begin{aligned} -\frac{1}{n}\log p(X_1,\ldots,X_n) &= -\frac{1}{n}\sum_i \log p(X_i) \\ &= \frac{1}{n}(Z_1+\ldots+Z_n) \\ &= \overline{Z}. \end{aligned} \tag{8}$$

So, by the weak law of large numbers, for $\epsilon, \delta > 0$ and sufficiently large $n$:

$$\Pr(|\overline{Z} - < Z > | \geq \delta) < \epsilon \tag{9}$$

i.e.,

$$\Pr(| - \frac{1}{n}\log p(X_1,\ldots,X_n) - H(X)| \geq \delta) < \epsilon \tag{10}$$

or equivalently,

$$\Pr(| - \frac{1}{n}\log p(X_1,\ldots,X_n) - H(X)| < \delta) \geq 1-\epsilon \tag{11}$$

and hence, with probability greater than or equal to $1-\epsilon$:

$$-n(H(X)+\delta) < \log p(X_1,\ldots,X_n) < -n(H(X)-\delta). \tag{12}$$

A '$\delta$-typical $n$-length sequence' $(x_1,\ldots,x_n) \in \mathcal{X}^n$ of values of the random variables $X_1,\ldots,X_n$ is defined as a sequence of symbols of $\mathcal{X}$ satisfying:

$$2^{-n(H(X)+\delta)} \leq p(x_1,\ldots,x_n) \leq 2^{-n(H(X)-\delta)}. \tag{13}$$

Denote the set of $\delta$-typical $n$-length sequences by $T_\delta^{(n)}$ and the number of sequences in $T_\delta^{(n)}$ by $|T_\delta^{(n)}|$. Then, for sufficiently large n,

$$\Pr(\{X_1,\ldots,X_n\} \in T_\delta^{(n)}) \geq 1-\epsilon; \tag{14}$$

and it can be shown that

$$(1-\epsilon)2^{n(H(X)-\delta)} \leq |T_\delta^{(n)}| \leq 2^{n(H(X)+\delta)}. \tag{15}$$

So, roughly, $T^{(n)}$ contains $2^{nH}$ equiprobable sequences, each having a probability of $2^{-nH}$.

Shannon's source coding theorem applies the above result about typical sequences to show that the compression rate of $H(X)$ bits per symbol produced by a source of independent and identically distributed random variables is optimal. The source produces $n$-length sequences of symbols $x_1, x_2, \ldots, x_n$ with probability $p(x_1, x_2, \ldots, x_n) = p(x_1)p(x_2)\ldots p(x_n)$, where each symbol is chosen from an alphabet $\mathcal{X}$. If there are $k$ symbols in $\mathcal{X}$, these $n$-sequences can be represented as sequences of $n \log k$ bits. Suppose there is a 'block coding' compression scheme that encodes each 'block' or

$n$-length sequence (for sufficiently large $n$) as a shorter sequence of $nR$ bits, where $0 \leq R \leq \log k$. Suppose also that the receiver has a decompression scheme for decoding sequences of $nR$ bits into sequences of $n$ symbols. Then one speaks of a compression/decompression scheme of *rate* $R$.

The source coding theorem states that

> if the Shannon entropy of a source is $H(X)$, then there exists a reliable compression/decompression scheme of rate $R$ if and only if $R \geq H(X)$, where a scheme is said to be reliable if it reproduces the original sequence with a probability that tends to 1 as $n \to \infty$.

For reliable communication, we want the compression and decompression of a sequence of symbols to yield the original sequence, but in general there will be a certain probability, $q(x_1, \ldots, x_n)$, of decoding a given sequence of $nR$ encoded bits received by the receiver as the original $n$-sequence produced by the source. The average *fidelity*[3] of a compression/decompression scheme for $n$-length blocks is defined as:

$$F_n = \sum_{\text{all } n\text{-sequences}} p(x_1, \ldots, x_n) q(x_1, \ldots, x_n) \tag{16}$$

If all the probabilities $q(x_1, \ldots, x_n)$ are 1, $F_n = 1$; otherwise $F_n < 1$. In terms of the fidelity as a measure of reliability of correct decoding, the source coding theorem states that

> for any $\epsilon, \delta > 0$: (i) there exists a compression/decompression scheme using $H(X) + \delta$ bits per symbol for $n$-length sequences produced by the source that can be decompressed by the receiver with a fidelity $F_n > 1 - \epsilon$, for sufficiently large $n$, and (ii) any compression/decompression scheme using $H(X) - \delta$ bits per symbol for $n$-length sequences will have a fidelity $F_n < \epsilon$, for sufficiently large $n$.

As a simple example of compression, consider an information source that produces sequences of symbols from a 4-symbol alphabet $a_1, a_2, a_3, a_4$ with probabilities 1/2, 1/4, 1/8, 1/8. Each symbol can be represented by a distinct 2-digit binary number:

$$\begin{aligned} a_1 &: 00 \\ a_2 &: 01 \\ a_3 &: 10 \\ a_4 &: 11 \end{aligned}$$

so without compression we need two bits per symbol of storage space to store the output of the source. The Shannon entropy of the source is $H(X) = -\frac{1}{2}\log\frac{1}{2} - \frac{1}{4}\log\frac{1}{4} - \frac{1}{8}\log\frac{1}{8} - \frac{1}{8}\log\frac{1}{8} = \frac{7}{4}$. Shannon's source coding theorem tells us that there is a compression scheme that uses an average of 7/4 bits per symbol rather than two

[3] Note that this definition of fidelity is different from the definition proposed by Nielsen and Chuang [2000, p.400] for the fidelity between two probability distributions $\{p_x\}$ and $\{q_x\}$ as a 'distance measure' between the distributions. They define $F_{NC}(p_x, q_x) := \sum_x \sqrt{p_x q_x}$, so $F_{NC}(p_x, q_x) = 1$ if $p_x = q_x$.

bits per symbol, and that such a compression scheme is optimal. The optimal scheme is provided by the following encoding:

$$\begin{aligned} a_1 &: 0 \\ a_2 &: 10 \\ a_3 &: 110 \\ a_4 &: 111 \end{aligned}$$

for which the *average* length of a compressed sequence is: $\frac{1}{2}\cdot 1+\frac{1}{4}\cdot 2+\frac{1}{8}\cdot 3+\frac{1}{8}\cdot 3 = \frac{7}{4}$ bits per symbol.

The significance of Shannon's source coding theorem lies is showing that there is an optimal or most efficient way of compressing messages produced by a source (assuming a certain idealization) in such a way that they can be reliably reconstructed by a receiver. Since a message is abstracted as a sequence of distinguishable symbols produced by a stochastic source, the only relevant feature of a message with respect to reliable compression and decompression is the sequence of probabilities associated with the individual symbols: the nature of the physical systems embodying the representation of the message through their states is irrelevant to this notion of compression (provided only that the states are reliably distinguishable), as is the content or meaning of the message. The Shannon entropy $H(X)$ is a measure of the minimal physical resources, in terms of the average number of bits per symbol, that are necessary and sufficient to reliably store the output of a source of messages. In this sense, it is a measure of the amount of information per symbol produced by an information source.

The essential notion underlying Shannon's measure of information is compressibility: information as a physical resource is something that can be compressed, and the amount of information produced by an information source is measured by its optimal compressibility.

## 2.2 Conditional Entropy, Mutual Information, Channel Capacity

The analysis so far assumes a noiseless channel between the source and the receiver. I turn now to a brief sketch of some concepts relevant to a noisy channel, and a statement of Shannon's noisy channel coding theorem.

An information channel maps inputs consisting of values of a random variable $X$ onto outputs consisting of values of a random variable $Y$, and the map will generally not be 1-1 if the channel is noisy. Consider the conditional probabilities $p(y|x)$ of obtaining an output value $y$ for a given input value $x$, for all $x, y$. From the probabilities $p(x)$ we can calculate $p(y)$ as:

$$p(y) = \sum_x p(y|x)p(x)$$

and we can also calculate $p(x|y)$ by Bayes' rule from the probabilities $p(y|x)$ and $p(x)$, for all $x, y$, and hence the Shannon entropy of the conditional distribution $p(x|y)$, for all $x$ and a fixed $y$, denoted by $H(X|Y = y)$.

The quantity

$$H(X|Y) = \sum_y p(y) H(X|Y=y) \tag{17}$$

is known as the *conditional entropy*. It is the expected value of $H(X|Y=y)$ for all $y$. If we think of $H(X)$, the entropy of the distribution $\{p(x) : x \in \mathcal{X}\}$, as a measure of the uncertainty of the $X$-value, then $H(X|Y=y)$ is a measure of the uncertainty of the $X$-value, given the $Y$-value $y$, and $H(X|Y)$ is a measure of the average uncertainty of the $X$-value, given a $Y$-value.

Putting it differently, the number of input sequences of length $n$ that are consistent with a given output sequence (as $n \to \infty$) is $2^{nH(X|Y)}$, i.e., $H(X|Y)$ is the number of bits per symbol of additional information needed, on average, to identify an input $X$-sequence from a given $Y$-sequence. This follows because there are $2^{nH(X,Y)}$ typical sequences of pairs $(x, y)$, where the *joint entropy* $H(X,Y)$ is calculated from the joint probability $p(x,y)$. So there are

$$\frac{2^{nH(X,Y)}}{2^{nH(Y)}} = 2^{n(H(X,Y)-H(Y))} = 2^{nH(X|Y)} \tag{18}$$

typical $X$-sequences associated with a given $Y$-sequence.

The 'chain rule' equality

$$H(X,Y) = H(X) + H(Y|X) = H(Y) + H(X|Y) = H(Y,X) \tag{19}$$

follows immediately from the logarithmic definitions of the quantities:

$$\begin{aligned} H(X\!:\!Y) &:= -\sum_{x,y} p(x,y) \log p(x,y) \\ &= -\sum_{x,y} p(x)p(y|x) \log\left(p(x)p(y|x)\right) \\ &= -\sum_{x,y} p(x)p(y|x) \log p(x) - \sum_{x,y} p(x)p(y|x) \log p(y|x) \\ &= -\sum_x p(x) \log p(x) + \sum_x p(x) \left(-\sum_y p(y|x) \log p(y|x)\right) \\ &= H(X) + H(Y|X) \end{aligned} \tag{20}$$

Note that $H(X|Y) \neq H(Y|X)$.

The *mutual information* measures the average amount of information gained about $X$ by ascertaining a $Y$-value, i.e., the amount of information one random variable contains about another, or the reduction in uncertainty of one random variable obtained by measuring another.

Mutual information can be defined in terms of the concept of *relative entropy*, which is a measure of something like the distance between two probability distributions (although it is not a true metric, since it is not symmetric and does not satisfy the triangle inequality). The relative entropy between distributions $p(x)$ and $q(x)$ is defined as:

$$D(p \parallel q) = \sum_{x \in \mathcal{A}} p(x) \log \frac{p(x)}{q(x)}. \tag{21}$$

The mutual information can now be defined as:

$$\begin{aligned} H(X{:}Y) &= D(p(x,y) \parallel p(x)p(y)) \\ &= \sum_x \sum_y p(x,y) \log \frac{p(x,y)}{p(x)p(y)}. \end{aligned} \tag{22}$$

It follows that

$$H(X{:}Y) = H(X) - H(X|Y) = H(Y) - H(Y|X), \tag{23}$$

i.e., the mutual information of two random variables represents the average information gain about one random variable obtained by measuring the other: the difference between the initial uncertainty of one of the random variables, and the average residual uncertainty of that random variable after ascertaining the value of the other random variable. Also, since $H(X,Y) = H(X) + H(Y|X)$, it follows that

$$H(X{:}Y) = H(X) + H(Y) - H(X,Y); \tag{24}$$

i.e., the mutual information of two random variables is a measure of how much information they have in common: the sum of the information content of the two random variables, as measured by the Shannon entropy (in which joint information is counted twice), minus their joint information. Note that $H(X : X) = H(X)$, as we would expect.

For a noisy channel, if $X$ represents the input to the channel and $Y$ represents the output of the channel, $H(X{:}Y)$ represents the average amount of information gained about the input $X$ by ascertaining the value of the output $Y$. The *capacity* of a channel, $C$, is defined as the supremum of $H(X{:}Y)$ over all input distributions.

Shannon's noisy channel coding theorem shows, perhaps surprisingly, that up to $C$ bits of information can be sent through a noisy channel with arbitrary low error rate. That is,

> there exists an optimal coding for an information source with entropy $H \leq C$ such that $n$-length sequences produced by the source can be transmitted faithfully over the channel: the error rate tends to zero as $n \to \infty$. The probability of error tends to 1 if we attempt to transmit more than $C$ bits through the channel.

This means that there are two ways of improving the transmission rate over a noisy channel such as a telephone cable. We can improve the channel capacity by replacing the cable with a faster one, or we can improve the information processing (the data compression).

# 3 Quantum Information

The physical notion of information, discussed in §2, is profoundly transformed by the transition from classical mechanics to quantum mechanics. The aim of this section

is to bring out the nature of this transformation. In §3.1, I develop some core concepts of quantum mechanics relevant to quantum information: entangled states, the Schmidt decomposition, the density operator formalism for the representation of pure and mixed states, the 'purification' of mixed states, generalized quantum measurements in terms of positive operator valued measures (POVMs), and the evolution of open systems represented by quantum operations. I assume throughout Hilbert spaces of finite dimension (and so avoid all the technicalities of functional analysis required for the treatment of infinite-dimensional Hilbert spaces). In fact, there is no loss of generality here, since both classical and quantum information sources are considered to produce messages consisting of sequences of symbols from some finite alphabet, which we represent in terms of a finite set of classical or quantum states. Moreover, all the conceptual issues relevant to the difference between classical and quantum information show up in finite-dimensional Hilbert spaces. In §3.2, I introduce von Neumann's generalization of the Shannon entropy and related notions for quantum information. In §3.3 and §3.4, I discuss some salient features that distinguish quantum information from classical information: §3.3 deals with the limitations on copying quantum information imposed by the 'no cloning' theorem, and §3.4 deals with the limited accessibility of quantum information defined by the Holevo bound. Finally, in §3.5 I show how the notion of compressibility applies to quantum information, and I outline Schumacher's generalization of Shannon's source coding theorem for quantum information, noting a distinction between 'visible' and 'blind' compression applicable to quantum information.

## 3.1 Some Relevant Quantum Mechanics

### 3.1.1 Entangled States

Consider a quantum system $Q$ which is part of a compound system $QE$; $E$ for 'environment,' although $E$ could be any quantum system of which $Q$ is a subsystem. Pure states of $QE$ are represented as rays or unit vectors in a tensor product Hilbert space $\mathcal{H}^Q \otimes \mathcal{H}^E$. A general pure state of $QE$ is a state of the form:

$$|\Psi\rangle = \sum c_{ij}|q_i\rangle|e_j\rangle \tag{25}$$

where $|q_i\rangle \in \mathcal{H}^Q$ is a complete set of orthonormal states (a basis) in $\mathcal{H}^Q$ and $|e_j\rangle \in \mathcal{H}^E$ is a basis in $\mathcal{H}^E$. If the coefficients $c_{ij}$ are such that $|\Psi\rangle$ cannot be expressed as a product state $|Q\rangle|E\rangle$, then $|\Psi\rangle$ is called an *entangled state*.

For any state $|\Psi\rangle$ of $QE$, there exist orthonormal bases $|i\rangle \in \mathcal{H}^Q$, $|j\rangle \in \mathcal{H}^E$ such that $|\Psi\rangle$ can be expressed in a biorthogonal correlated form as:

$$|\Psi\rangle = \sum_i \sqrt{p_i}|i\rangle|i\rangle \tag{26}$$

where the coefficients $\sqrt{p_i}$ are real and non-negative, and $\sum p_i = 1$. This representation is referred to as the *Schmidt decomposition*. The Schmidt decomposition is unique if and only if the $p_i$ are all distinct.

An example is the biorthogonal EPR state[4]

$$|\Psi\rangle = (|0\rangle|1\rangle - |1\rangle|0\rangle)/\sqrt{2}; \tag{27}$$

say, the singlet state of two spin-1/2 particles (the Schmidt form with positive coefficients is obtained by asborbing the relative phases in the definition of the basis vectors). In the singlet state, $|0\rangle$ and $|1\rangle$ can be taken as representing the two eigenstates of spin in the $z$-direction, but since the state is symmetric, $|\Psi\rangle$ retains the same form for spin in any direction. The EPR argument exploits the fact that spin measurements in the same direction on the two particles, which could be arbitrarily far apart, will yield outcomes that are perfectly anti-correlated for any spin direction. Bell's counterargument exploits the fact that when the spin is measured on one particle in a direction $\theta_1$ to the $z$-axis, but on the other particle in a direction $\theta_2$ to the $z$-axis, the probability of finding the same outcome for both particles (both 1 or both 0) is $\frac{1}{2}\sin^2(\theta_1 - \theta_2)$. It follows that the outcomes are perfectly correlated when $\theta_1 - \theta_2 = \pi$ and that 3/4 of the outcomes are the same when $\theta_1 - \theta_2 = 2\pi/3$. On the other hand, from Bell's inequality, derived under Einstein's realist assumptions of separability and locality, we see that the correlation for $\theta_1 - \theta_2 = 2\pi/3$ cannot exceed 2/3. See Dickson (this vol., ch. 4) for further discussion.

This means that the dynamical evolution of a quantum system can result in a state representing correlational information that no classical computer can simulate. That is, no classical computer can be programmed to perform the following task: for any pair of input angles, $\theta_1, \theta_2$, at different locations, output a pair of values (0 or 1) for these locations such that the values are perfectly correlated when $\theta_1 - \theta_2 = \pi$, perfectly anti-correlated when $\theta_1 = \theta_2$, and 75% correlated when $\theta_1 - \theta_2 = 2\pi/3$, where the response time between being given the input and producing the output in each case is less than the time taken by light to travel between the two locations.

Notice that the four states:

$$|00\rangle = \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|1\rangle) \tag{28}$$

$$|01\rangle = \frac{1}{\sqrt{2}}(|0\rangle|1\rangle + |1\rangle|0\rangle) \tag{29}$$

$$|10\rangle = \frac{1}{\sqrt{2}}(|0\rangle|0\rangle - |1\rangle|1\rangle) \tag{30}$$

$$|11\rangle = \frac{1}{\sqrt{2}}(|0\rangle|1\rangle - |1\rangle|0\rangle) \tag{31}$$

form an orthonormal basis, called the Bell basis, in the 2 x 2-dimensional Hilbert space. Any Bell state can be transformed into any other Bell state by a local unitary transformation, $X, Y$, or $Z$, where $X, Y, Z$ are the Pauli spin matrices:

$$X = \sigma_x = |0\rangle\langle 1| + |1\rangle\langle 0| = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \tag{32}$$

[4]Einstein, Podolsky and Rosen considered a more complicated state entangled over position and momentum values. The spin example is due to Bohm [1951, pp, 611–623].

$$Y = \sigma_y = i|0\rangle\langle 1| - i|1\rangle\langle 0| = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \tag{33}$$

$$Z = \sigma_z = |0\rangle\langle 0| + |1\rangle\langle 1| = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{34}$$

For example:

$$X \otimes I \cdot \frac{1}{\sqrt{2}}(|0\rangle\langle 1| - |1\rangle|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle\langle 0| - |1\rangle|1\rangle. \tag{35}$$

If $QE$ is a closed system in an entangled pure state represented by

$$|\Psi\rangle = \sum_i \sqrt{p_i}|i\rangle|i\rangle \tag{36}$$

in the Schmidt decomposition, the expected value of any $Q$-observable $A$ on $\mathcal{H}^Q$ can be computed as:

$$\begin{aligned} \langle A\rangle &= \mathrm{Tr}(|\Psi\rangle\langle\Psi|A \otimes I) \\ &= \mathrm{Tr}_Q(\mathrm{Tr}_E(|\Psi\rangle\langle\Psi|A)) \\ &= \mathrm{Tr}_Q(\sum_i p_i|i\rangle\langle i|A) \\ &= \mathrm{Tr}_Q(\rho A) \end{aligned} \tag{37}$$

where $\mathrm{Tr}_Q() = \sum_q \langle q_i| \cdot |q_i\rangle$, for any orthonormal basis in $\mathcal{H}^Q$, is the partial trace over $\mathcal{H}^Q$, $\mathrm{Tr}_E()$ is the partial trace over $\mathcal{H}^E$, and $\rho = \sum_i p_i|i\rangle\langle i| \in \mathcal{H}^Q$ is the *reduced density operator* of the open system $Q$, a positive operator with unit trace. Since the density operator $\rho$ yields the statistics of all $Q$-observables via Eq. (37), $\rho$ is taken as representing the quantum state of the system $Q$.

If $QE$ is an entangled pure state, then the open system $Q$ is in a *mixed state* $\rho$, i.e., $\rho \neq \rho^2$; for pure states, $\rho$ is a projection operator onto a ray and $\rho = \rho^2$. A mixed state represented by a density operator $\rho = \sum \rho_i|i\rangle\langle i|$ can be regarded as a mixture of pure states $|i\rangle$ prepared with prior probabilities $p_i$, but this representation is not unique—not even if the states combined in the mixture are orthogonal. For example, the equal-weight mixture of orthonormal states $|0\rangle, |1\rangle$ in a 2-dimensional Hilbert space $\mathcal{H}_2$ has precisely the same statistical properties, and hence the same density operator $\rho = I/2$, as the equal weight mixture of any pair of orthonormal states, e.g., the states $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle), \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$, or the equal-weight mixture of nonorthogonal states $|0\rangle, \frac{1}{2}|0\rangle + \frac{\sqrt{3}}{2}|1\rangle, \frac{1}{2}|0\rangle - \frac{\sqrt{3}}{2}|1\rangle$ $120°$ degrees apart, or the uniform continuous distribution over all possible states in $\mathcal{H}_2$.

More generally, for any basis of orthonormal states $|e_i\rangle \in \mathcal{H}^E$, the entangled state $|\Psi\rangle$ can be expressed as:

$$|\Psi\rangle = \sum_{ij} c_{ij}|q_i\rangle|e_j\rangle = \sum_j \sqrt{w_j}|r_j\rangle|e_j\rangle \tag{38}$$

where the normalized states $|r_j\rangle = \sum_i \frac{c_{ij}}{\sqrt{w_j}}|q_i\rangle$ are *relative states* to the $|e_j\rangle$ ($\sqrt{w_j} = \sum_j |c_{ij}|^2$). Note that the states $|r_j\rangle$ are not in general orthogonal. Since the $|e_j\rangle$ are orthogonal, we can express the density operator representing the state of $Q$ as:

$$\rho = \sum_i w_i |r_i\rangle\langle r_i|. \tag{39}$$

In effect, a measurement of an $E$-observable with eigenstates $|e_i\rangle$ will leave the composite system $QE$ in one of the states $|r_i\rangle|e_i\rangle$ with probability $w_i$, and a measurement of an $E$-observable with eigenstates $|i\rangle$ (the orthogonal states of the Schmidt decomposition in (36) above) will leave the system $QE$ in one of the states $|i\rangle|i\rangle$ with probability $p_i$. Since $Q$ and $E$ could be widely separated from each other in space, no measurement at $E$ could affect the statistics of any $Q$-observable; or else measurements at $E$ would allow superluminal signaling between $Q$ and $E$. It follows that the mixed state $\rho$ can be realized as a mixture of orthogonal states $|i\rangle$ (the eigenstates of $\rho$) with weights $p_i$, or as a mixture of non-orthogonal relative states $|r_j\rangle$ with weights $w_j$ in infinitely many ways, depending on the choice of basis in $\mathcal{H}^E$:

$$\rho = \sum_i p_i |i\rangle\langle i| = \sum_j w_j |r_j\rangle\langle r_j| \tag{40}$$

and all these different mixtures with the same density operator $\rho$ must be physically indistinguishable.

Note that any mixed state density operator $\rho \in \mathcal{H}^Q$ can be 'purified' by adding a suitable ancilla system $E$, in the sense that $\rho$ is the partial trace of a pure state $|\Psi\rangle \in \mathcal{H}^Q \otimes \mathcal{H}^E$ over $\mathcal{H}^E$. A purification of a mixed state is, clearly, not unique, but depends on the choice of $|\Psi\rangle$ in $\mathcal{H}^E$. The Hughston-Jozsa-Wootters theorem [Hughston *et al.*, 1993] shows that for *any* mixture of pure states $|r_i\rangle$ with weights $w_i$, where $\rho = \sum_j w_j |r_j\rangle\langle r_j|$, there is a purification of $\rho$ and a suitable measurement on the system $E$ that will leave $Q$ in the mixture $\rho$. So an observer at $E$ can remotely prepare $Q$ in any mixture that corresponds to the density operator $\rho$ (and of course all these different mixtures are physically indistinguishable). Similar results were proved earlier by Schrödinger [1936], Jaynes [1957] and Gisin [1989]. See Halvorson [2004] for a generalization to hyperfinite von Neuman algebras.

### 3.1.2 Measurement

A standard von Neumann 'yes-no' measurement is associated with a projection operator; so a standard observable is represented in the spectral representation as a sum of projection operators, with coefficients representing the eigenvalues of the observable. Such a measurement is the quantum analogue of the measurement of a property of a system in classical physics. Classically, we think of a property of a system as being associated with a subset in the state space (phase space) of the system, and determining whether the system has the property amounts to determining whether the state of the system lies in the corresponding subset. In quantum mechanics, the counterpart of a subset in phase space is a closed linear subspace in Hilbert space. Just as the different possible values of an observable (dynamical quantity) of a classical system correspond

to the subsets in a mutually exclusive and collectively exhaustive set of subsets covering the classical state space, so the different values of a quantum observable correspond to the subspaces in a mutually exclusive (i.e., orthogonal) and collectively exhaustive set of subspaces spanning the quantum state space. (For further discussion, see Dickson, this vol., ch. 4, [Mackey, 1963], and [Bub, 1997].)

In quantum mechanics, and especially in the theory of quantum information (where any read-out of the quantum information encoded in a quantum state requires a quantum measurement), it is useful to consider a more general class of measurements than the projective measurements associated with the determination of the value of an observable. It is common to speak of generalized measurements and generalized observables. But in fact this terminology is more misleading than illuminating, because a generalized measurement is not a procedure that reveals whether or not a quantum system has some sort of generalized property. Rather, the point of the generalization is to exploit the difference between quantum and classical states for new possibilities in the representation and manipulation of information.

To clarify the idea, I will follow the excellent discussion by Nielsen and Chuang [2000, §2.2.3–2.2.6]. A quantum measurement can be characterized, completely generally, as a certain sort of interaction between two quantum systems, $Q$ (the measured system) and $M$ (the measuring system). We suppose that $Q$ is initially in a state $|\psi\rangle$ and that $M$ is initially in some standard state $|0\rangle$, where $|m\rangle$ is an orthonormal basis of 'pointer' eigenstates in $\mathcal{H}^M$. The interaction is defined by a unitary transformation $U$ on the Hilbert space $\mathcal{H}^Q \otimes \mathcal{H}^M$ that yields the transition:

$$|\psi\rangle|0\rangle \xrightarrow{U} \sum_m M_m|\psi\rangle|m\rangle \tag{41}$$

where $\{M_m\}$ is a set of linear operators (the Kraus operators) defined on $\mathcal{H}^Q$ satisfying the *completeness condition*:

$$\sum_m M_m^\dagger M_m = I. \tag{42}$$

(The symbol † denotes the adjoint or Hermitian conjugate.) The completeness condition guarantees that this evolution is unitary, because it guarantees that $U$ preserves inner products, i.e.

$$\begin{aligned}
\langle\phi|\langle 0|U^\dagger U|\psi\rangle|0\rangle &= \sum_{m,m'}\langle m|\langle\phi|M_m^\dagger M_{m'}|\psi\rangle|m'\rangle \\
&= \sum_m \langle\phi|M^\dagger M|\psi\rangle \\
&= \langle\phi|\psi\rangle
\end{aligned} \tag{43}$$

from which it follows that $U$, defined as above by Eq. (41) for any product state $|\psi\rangle|0\rangle$ (for any $|\psi\rangle \in \mathcal{H}^Q$) can be extended to a unitary operator on the Hilbert space $\mathcal{H}^Q \otimes \mathcal{H}^M$. Accordingly, any set of linear operators $\{M_m\}$ defined on the Hilbert space of the system $Q$ satisfying the completeness condition defines a measurement in this general sense, with the index $m$ labeling the possible outcomes of the measurement, and any such set is referred to as a set of measurement operators.

If we now perform a standard projective measurement on $M$ to determine the value $m$ of the pointer observable, defined by the projection operator

$$P_m = I_Q \otimes |m\rangle\langle m|$$

then the probability of obtaining the outcome $m$ is, by (37)[5]:

$$\begin{aligned} p(m) &= \langle 0|\langle\psi|U^\dagger P_m U|\psi\rangle|0\rangle \\ &= \sum_{m'm''} \langle m'|\langle\psi|M^\dagger_{m'}(I_Q \otimes |m\rangle\langle m|)M_{m''}|\psi\rangle|m''\rangle \\ &= \sum_{m'm''} \langle\psi|M^\dagger_{m'}\langle m'|m\rangle\langle m|m''\rangle M_{m''}|\psi\rangle \\ &= \langle\psi|M^\dagger_m M_m|\psi\rangle; \end{aligned} \tag{44}$$

and, more generally, if the initial state of $Q$ is a mixed state $\rho$, then

$$p(m) = \mathrm{Tr}_Q(M\rho M^\dagger). \tag{45}$$

The final state of $QM$ after the projective measurement on $M$ yielding the outcome $m$ is:

$$\frac{P_m U|\psi\rangle|0\rangle}{\sqrt{\langle\psi|U^\dagger P U|\psi\rangle}} = \frac{M_m|\psi\rangle|m\rangle}{\sqrt{\langle\psi|M^\dagger_m M_m|\psi\rangle}}. \tag{46}$$

So the final state of $M$ is $|m\rangle$ and the final state of $Q$ is:

$$\frac{M_m|\psi\rangle}{\sqrt{\langle\psi|M^\dagger_m M_m|\psi\rangle}};$$

and, more generally, if the initial state of $Q$ is a mixed state $\rho$, then the final state of $Q$ is:

$$\frac{M_m\rho M^\dagger_m}{\mathrm{Tr}_Q(M_m\rho M^\dagger_m)}.$$

Note that this general notion of measurement covers the case of standard projective measurements. In this case $\{M_m\} = \{P_m\}$, where $\{P_m\}$ is the set of projection operators defined by the spectral measure of a standard quantum observable represented by a self-adjoint operator. It also covers the measurement of 'generalized observables' associated with positive operator valued measures (POVMs). Let

$$E_m = M^\dagger_m M_m \tag{47}$$

then the set $\{E_m\}$ defines a set of positive operators ('effects') such that

$$\sum E_m = I \tag{48}$$

[5] The expected value of a projection operator, which is an idempotent observable with eigenvalues 0 and 1, is equal to the probability of obtaining the eigenvalue 1. Here the eigenvalue 1 corresponds to the outcome $m$.

A POVM can be regarded as a generalization of a projection valued measure (PVM), in the sense that Eq. (48) defines a 'resolution of the identity' without requiring the PVM orthogonality condition:

$$P_m P_{m'} = \delta_{mm'} P_m. \tag{49}$$

Note that for a POVM:

$$p(m) = \langle\psi|E_m|\psi\rangle. \tag{50}$$

Given a set of positive operators $\{E_m\}$ such that $\sum E_m = I$, measurement operators $M_m$ can be defined via

$$M_m = U\sqrt{E_m}, \tag{51}$$

where $U$ is a unitary operator, from which it follows that

$$\sum_m M_m^\dagger M_m = \sum E_m = I \tag{52}$$

As a special case, of course, we can take $U = 1$ and $M_m = \sqrt{E_m}$. Conversely, given a set of measurement operators $\{M_m\}$, there exist unitary operators $U_m$ such that $M_m = U_m\sqrt{E_m}$, where $\{E_m\}$ is a POVM. (This follows immediately from [Nielsen and Chuang, 2000, Theorem 2.3, p. 78]; see [Nielsen and Chuang, 2000, Exercise 2.63, p. 92].)

Except for the standard case of projective measurements, one might wonder why it might be useful to single out such unitary transformations, and why in the general case such a process should be called a *measurement* of $Q$. The following example, taken from [Nielsen and Chuang, 2000, p. 92], is illuminating. Suppose we know that a system with a 2-dimensional Hilbert space is in one of two nonorthogonal states:

$$\begin{aligned} |\psi_1\rangle &= |0\rangle \\ |\psi_2\rangle &= \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \end{aligned}$$

It is impossible to *reliably* distinguish these states by a quantum measurement, even in the above generalized sense. Here 'reliably' means that the state is identified correctly with zero probability of error.

To see this, suppose there is such a measurement, defined by two measurement operators $M_1, M_2$ satisfying the completeness condition. Then we require

$$p(1) = \langle\psi_1|M_1^\dagger M_1|\psi_1\rangle = 1, \tag{53}$$

to represent reliability if the state is $|\psi_1\rangle$; and

$$p(2) = \langle\psi_2|M_2^\dagger M_2|\psi_2\rangle = 1 \tag{54}$$

to represent reliability if the state is $|\psi_2\rangle$. By the completeness condition we must have

$$\langle\psi_1|M_1^\dagger M_1 + M_2^\dagger M_2|\psi_1\rangle = 1 \tag{55}$$

from which it follows that $\langle\psi_1|M_2^\dagger M_2|\psi_1\rangle = 0$, i.e., $M_2|\psi_1\rangle = M_2|0\rangle = 0$. Hence

$$M_2|\psi_2\rangle = M_2\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) = \frac{1}{\sqrt{2}}M_2|1\rangle \tag{56}$$

and so

$$p(2) = \langle\psi_2|M_2^\dagger M_2|\psi_2\rangle = \frac{1}{2}\langle 1|M_2^\dagger M_2|1\rangle. \tag{57}$$

But by the completeness condition we also have

$$\langle 1|M_2^\dagger M_2|1\rangle \leq \langle 1|M_1^\dagger M_1 + M_2^\dagger M_2|1\rangle = \langle 1|1\rangle = 1 \tag{58}$$

from which it follows that

$$p(2) \leq \frac{1}{2} \tag{59}$$

which contradicts Eq. (54).

However, it is possible to perform a measurement in the generalized sense, with *three* possible outcomes, that will allow us to correctly identify the state some of the time, i.e., for two of the possible outcomes, while nothing about the identity of the state can be inferred from the third outcome.

Here's how: The three operators

$$\begin{aligned} E_1 &= \frac{\sqrt{2}}{1+\sqrt{2}}\frac{(|0\rangle - |1\rangle)(\langle 0| - \langle 1|)}{2} \\ E_2 &= \frac{\sqrt{2}}{1+\sqrt{2}}|1\rangle\langle 1| \\ E_3 &= I - E_1 - E_2 \end{aligned} \tag{60}$$

are all positive operators and $E_1 + E_2 + E_3 = I$, so they define a POVM. In fact, $E_1, E_2, E_3$ are each multiples of projection operators onto the states

$$\begin{aligned} |\phi_1\rangle &= |\psi_2\rangle^\perp \\ |\phi_2\rangle &= |\psi_1\rangle^\perp \\ |\phi_3\rangle &= \frac{(1+\sqrt{2})|0\rangle + |1\rangle}{\sqrt{2\sqrt{2}(1+\sqrt{2})}} \end{aligned} \tag{61}$$

with coefficients $\frac{\sqrt{2}}{1+\sqrt{2}}, \frac{\sqrt{2}}{1+\sqrt{2}}, \frac{1}{1+\sqrt{2}}$ respectively. The measurement involves a system $M$ with three orthogonal pointer states $|1\rangle, |2\rangle, |3\rangle$. The appropriate unitary interaction $U$ results in the transition, for an input state $|\psi\rangle$:

$$|\psi\rangle|0\rangle \xrightarrow{U} \sum_m M_m|\psi\rangle|m\rangle \tag{62}$$

where $M_m = \sqrt{E_m}$.

If the input state is $|\psi_1\rangle = |0\rangle$, we have the transition:

$$\begin{aligned} |\psi_1\rangle|0\rangle &\xrightarrow{U} \sqrt{E}_1|0\rangle|1\rangle + \sqrt{E}_3|0\rangle|3\rangle \\ &= \alpha|\phi_1\rangle|1\rangle + \beta|\phi_3\rangle|3\rangle \end{aligned} \tag{63}$$

(because $\sqrt{E_2}|\psi_1\rangle = \sqrt{E_2}|0\rangle = 0$). And if the input state is $|\psi_2\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, we have the transition:

$$\begin{aligned} |\psi_2\rangle|0\rangle &\xrightarrow{U} \sqrt{E_2}\frac{|0\rangle + |1\rangle}{\sqrt{2}}|2\rangle + \sqrt{E_3}\frac{|0\rangle + |1\rangle}{\sqrt{2}}|3\rangle \\ &= \gamma|\phi_2\rangle|2\rangle + \delta|\phi_3\rangle|3\rangle \end{aligned} \tag{64}$$

(because $\sqrt{E_1}|\psi_2\rangle = \sqrt{E_1}\frac{|0\rangle+|1\rangle}{\sqrt{2}} = 0$), where $\alpha, \beta, \gamma, \delta$ are real numerical coefficients.

We see that a projective measurement of the pointer of $M$ that yields the outcome $m = 1$ indicates, with certainty, that the input state was $|\psi_1\rangle = |0\rangle$. In this case, the measurement leaves the system $Q$ in the state $|\phi_1\rangle$. A measurement outcome $m = 2$ indicates, with certainty, that the input state was $|\psi_2\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, and in this case the measurement leaves the system Q in the state $|\phi_2\rangle$. If the outcome is $m = 3$, the input state could have been either $|\psi_1\rangle = |0\rangle$ or $|\psi_2\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, and $Q$ is left in the state $|\phi_3\rangle$.

### 3.1.3 Quantum Operations

When a closed system $QE$ evolves under a unitary transformation, $Q$ can be shown to evolve under a *quantum operation*, i.e., a completely positive linear map:

$$\mathcal{E} : \rho \to \rho' \tag{65}$$

where

$$\mathcal{E}(\rho) = \mathrm{Tr}_E(U\rho \otimes \rho_E U^\dagger) \tag{66}$$

(See [Nielsen and Chuang, 2000, p. 356 ff].) The map $\mathcal{E}$ is linear (or convex-linear) in the sense that $\mathcal{E}(\sum_i p_i\rho_i) = \sum_i p_i\mathcal{E}(p_i)$, positive in the sense that $\mathcal{E}$ maps positive operators to positive operators, and completely positive in the sense that $\mathcal{E} \otimes I$ is a positive map on the extension of $\mathcal{H}^Q$ to a Hilbert space $\mathcal{H}^Q \otimes \mathcal{H}^E$, associated with the addition of any ancilla system $E$ to $Q$.

Every quantum operation (i.e., completely positive linear map) on a Hilbert space $\mathcal{H}^Q$ has a (non-unique) representation as a unitary evolution on an extended Hilbert space $\mathcal{H}^Q \otimes \mathcal{H}^E$, i.e.,

$$\mathcal{E}(\rho) = \mathrm{Tr}_E(U(\rho \otimes \rho_E)U^\dagger) \tag{67}$$

where $\rho_E$ is an appropriately chosen initial state of an ancilla system $E$ (which we can think of as the environment of $Q$). It turns out that it suffices to take $\rho_E$ as a pure state, i.e., $|0\rangle\langle 0|$, since a mixed state of $E$ can always be purified by enlarging the Hilbert space (i.e., adding a further ancilla system). So the evolution of a system $Q$ described by a quantum operation can always be modeled as the unitary evolution of a system $QE$, for an initial pure state of $E$.

Also, every quantum operation on a Hilbert space $\mathcal{H}^Q$ has a (non-unique) operator sum representation intrinsic to $\mathcal{H}^Q$:

$$\mathcal{E}(\rho) = \sum_i E_i\rho E_i^\dagger \tag{68}$$

where $E_i = \langle i|U|0\rangle$ for some orthonormal basis $\{|i\rangle\}$ of $E$. (See [Nielsen and Chuang, 2000, Theorem 8.1, p. 368].) If the operation is trace-preserving (or nonselective), then $\sum_i E_i^\dagger E_i = I$. For operations that are not trace-preserving (or selective), $\sum_i E_i^\dagger E_i \leq I$. This corresponds to the case where the outcome of a measurement on $QE$ is taken into account (selected) in the transition $\mathcal{E} \rightarrow \mathcal{E}(\rho)$.

If there is no interaction between $Q$ and $E$, then $\epsilon(\rho) = U_Q \rho U_Q^\dagger$, $U_Q U_Q^\dagger = I$, i.e., there is only one operator in the sum. In this case, $U = U_Q \otimes U_E$ and

$$\begin{aligned} \mathcal{E}(\rho) &= \mathrm{Tr}_E(U_Q \otimes U_E(\rho \otimes |0\rangle\langle 0|)U_Q^\dagger \otimes U_E^\dagger) & (69) \\ &= U_Q \rho U_Q^\dagger. & (70) \end{aligned}$$

So unitary evolution is a special case of the operator sum representation of a quantum operation and, of course, another special case is the transition $\mathcal{E} \rightarrow \mathcal{E}(\rho)$ that occurs in a quantum measurement process, where $E_i = M_i$. A trace-preserving operation corresponds to a non-selective measurement:

$$\mathcal{E}(\rho) = \sum_i M_i \rho M_i^\dagger; \tag{71}$$

while an operation that is not trace-preserving corresponds to a selective measurement, where the state 'collapses' onto the corresponding measurement outcome:

$$M_i \rho M_i^\dagger / \mathrm{Tr}(M_i \rho M_i^\dagger). \tag{72}$$

The operator sum representation applies to quantum operations between possibly different input and output Hilbert spaces, and characterizes the following general situation: a quantum system in an unknown initial state $\rho$ is allowed to interact unitarily with other systems prepared in standard states, after which some part of the composite system is discarded, leaving the final system in a state $\rho'$. The transition $\rho \rightarrow \rho'$ is defined by a quantum operation. So a quantum operation represents, quite generally, the unitary evolution of a closed quantum system, the nonunitary evolution of an open quantum system in interaction with its environment, and evolutions that result from a combination of unitary interactions and selective or nonselective measurements.

As we have seen, the creed of the Church of the Larger Hilbert Space is that every state can be made pure, every measurement can be made ideal, and every evolution can be made unitary – on a larger Hilbert space.[6]

## 3.2 Von Neumann Entropy

In this section, I define the von Neumann entropy of a mixture of quantum states (von Neumann's generalization of the Shannon entropy of a classical probability distribution characterizing a classical information source) and the corresponding notions of conditional entropy and mutual information.

Information in Shannon's sense is a quantifiable resource associated with the output of a (suitably idealized) stochastic source of symbolic states, where the physical

[6]The Creed originates with John Smolin. I owe this formulation to Ben Schumacher. See his *Lecture Notes on Quantum Information Theory* [1998].

nature of the systems embodying these states is irrelevant to the amount of classical information associated with the source. The quantity of information associated with a stochastic source is defined by its optimal compressibility, and this is given by the Shannon entropy. The fact that some feature of the output of a stochastic source can be optimally compressed is, ultimately, what justifies the attribution of a quantifiable resource to the source.

Information is represented physically in the states of physical systems. The essential difference between classical and quantum information arises because of the different distinguishability properties of classical and quantum states. As we will see below, only sets of orthogonal quantum states are reliably distinguishable (i.e., with zero probability of error), as are sets of different classical states (which are represented by disjoint singleton subsets in a phase space, and so are orthogonal as subsets of phase space in a sense analogous to orthogonal subspaces of a Hilbert space).

Classical information is that sort of information represented in a set of distinguishable states—states of classical systems, or orthogonal quantum states—and so can be regarded as a subcategory of quantum information, where the states may or may not be distinguishable. The idea behind quantum information is to extend Shannon's notion of compressibility to a stochastic source of quantum states, which may or may not be distinguishable. For this we need to define a suitable measure of information for probability distributions of quantum states—mixtures—as a generalization of the notion of Shannon entropy.

Consider a system $QE$ in an entangled state $|\Psi\rangle$. Then the subsystem $Q$ is in a mixed state $\rho$, which can always be expressed as:

$$\rho = \sum_i p_i |i\rangle\langle i| \tag{73}$$

where the $p_i$ are the eigenvalues of $\rho$ and the pure states $|i\rangle$ are orthonormal eigenstates of $\rho$. This is the spectral representation of $\rho$, and any density operator—a positive (hence Hermitian) operator—can be expressed in this way. The representation is unique if and only if the $p_i$ are all distinct. If some of the $p_i$ are equal, there is a unique representation of $\rho$ as a sum of projection operators with the distinct values of the $p_i$ as coefficients, but some of the projection operators will project onto multi-dimensional subspaces.

Since $\rho$ has unit trace, $\sum p_i = 1$, and so the spectral representation of $\rho$ represents a classical probability distribution of orthogonal, and hence distinguishable, pure states. If we measure a $Q$-observable with eigenstates $|i\rangle$, then the outcomes can be associated with the values of a random variable $X$, where $\Pr(X = i) = p_i$. Then

$$H(X) = -\sum p_i \log p_i \tag{74}$$

is the Shannon entropy of the probability distribution of measurement outcomes.

Now,

$$-\operatorname{Tr}(\rho \log \rho) = -\sum p_i \log p_i \tag{75}$$

(because the eigenvalues of $\rho \log \rho$ are $p_i \log p_i$ and the trace of an operator is the sum of the eigenvalues), so a natural generalization of Shannon entropy for any mixture of

quantum states with density operator $\rho$ is the *von Neumann entropy*[7]:

$$S := -\text{Tr}(\rho \log \rho) \tag{76}$$

which coincides with the Shannon entropy for measurements in the eigenbasis of $\rho$. For a completely mixed state $\rho = I/d$, where $\dim \mathcal{H}^Q = d$, the $d$ eigenvalues of $\rho$ are all equal to $1/d$ and $S = \log d$. This is the maximum value of $S$ in a $d$-dimensional Hilbert space. The von Neumann entropy $S$ is zero, the minimum value, if and only if $\rho$ is a pure state, where the eigenvalues of $\rho$ are 1 and 0. So $0 \leq S \leq \log d$, where $d$ is the dimension of $\mathcal{H}^Q$.

Recall that we can think of the Shannon entropy as a measure of the average amount of information gained by identifying the state produced by a known stochastic source. Alternatively, the Shannon entropy represents the optimal compressibility of the information produced by an information source. The von Neumann entropy does *not*, in general, represent the amount of information gained by identifying the quantum state produced by a stochastic source characterized as a mixed state, because nonorthogonal quantum states in a mixture cannot be reliably identified. However, as we will see in §3.5, the von Neumann entropy can be interpreted in terms of compressibility via Schumacher's source coding theorem Schumacher's source coding theorem for quantum information [Schumacher, 1995], a generalization of Shannon's source coding theorem for classical information. For an elementary two-state quantum system with a 2-dimensional Hilbert space considered as representing the output of an elementary quantum information source, $S = 1$ for an equal weight distribution over two orthogonal states (i.e., for the density operator $\rho = I/2$), so Schumacher takes the basic unit of quantum information as the 'qubit.' By analogy with the term 'bit,' the term 'qubit' refers to the basic unit of quantum information in terms of the von Neumann entropy, and to an elementary two-state quantum system considered as representing the possible outputs of an elementary quantum information source.

The difference between quantum information as measured by von Neumann entropy $S$ and classical information as measured by Shannon entropy $H$ can be brought out by considering the quantum notions of conditional entropy and mutual information (cf. §2.2), and in particular the peculiar feature of *inaccessibility* associated with quantum information.

For a composite system $AB$, conditional von Neumann entropy and mutual information are defined in terms of the joint entropy $S(AB) = -\text{Tr}(\rho^{AB} \log \rho^{AB})$ by analogy with the corresponding notions for Shannon entropy (cf. Eqs. (19), (23), (24)):

$$\begin{aligned} S(A|B) &= S(A,B) - S(B) && (77) \\ S(A\!:\!B) &= S(A) - S(A|B) && (78) \\ &= S(B) - S(B|A) && (79) \\ &= S(A) + S(B) - S(A,B) && (80) \end{aligned}$$

The joint entropy satisfies the subadditivity inequality:

$$S(A,B) \leq S(A) + S(B) \tag{81}$$

[7]Von Neumann first defined this quantity on the basis of a thermodynamic argument in [1955, p. 379].

with equality if and only if $A$ and $B$ are uncorrelated, i.e., $\rho^{AB} = \rho^A \otimes \rho^B$.

Now, $S(A|B)$ can be negative, while the conditional Shannon entropy is always positive or zero. Consider, for example, the entangled state $|\Psi\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$. Since $|\Psi\rangle$ is a pure state, $S(A,B) = 0$. But $S(A) = S(B) = 1$. So $S(A|B) = S(A,B) - S(A) = -1$. In fact, for a pure state $|\Psi\rangle$ of a composite system $AB$, $S(A|B) < 0$ if and only if $|\Psi\rangle$ is entangled.

For a composite system $AB$ in a product state $\rho \otimes \sigma$, it follows from the definition of joint entropy that:

$$S(A,B) = S(\rho \otimes \sigma) = S(\rho) + S(\sigma) = S(A) + S(B). \tag{82}$$

If $AB$ is in a pure state $|\Psi\rangle$, it follows from the Schmidt decomposition theorem that $|\Psi\rangle$ can be expressed as

$$|\Psi\rangle = \sum_i \sqrt{p_i}|i\rangle\langle i| \tag{83}$$

from which it follows that

$$\begin{array}{rclcl} \rho_A & = & \mathrm{Tr}_B(|\psi\rangle\langle\psi|) & = & \sum_i |i\rangle\langle i| \\ \rho_B & = & \mathrm{Tr}_A(|\psi\rangle\langle\psi|) & = & \sum_i |i\rangle\langle i|; \end{array} \tag{84}$$

and so:

$$S(A) = S(B) = -\sum_i p_i \log p_i. \tag{85}$$

Consider a mixed state prepared as a mixture of states $\rho_i$ with weights $p_i$. It can be shown that

$$S(\sum_i p_i\rho_i) \leq H(p_i) + \sum_i p_i S(\rho_i) \tag{86}$$

with equality if and only if the states $\rho_i$ have support on orthogonal subspaces (see [Nielsen and Chuang, 2000, Theorem 11.10, p. 518]). The entropy $H(p_i)$ is referred to as the *entropy of preparation* of the mixture $\rho$.

If the states $\rho_i$ are pure states, then $S(\rho) \leq H(p_i)$. For example, suppose $\mathcal{H}^Q$ is 2-dimensional and $p_1 = p_2 = 1/2$, then $H(p_i) = 1$. So if we had a classical information source producing the symbols 1 and 2 with equal probabilities, no compression of the information would be possible. However, if the symbols 1 and 2 are encoded as nonorthogonal quantum states $|r_1\rangle$ and $|r_2\rangle$, then $S(\rho) < 1$. As we will see in §3.5, according to Schumacher's source coding theorem, since $S(\rho) < 1$, quantum compression is possible, i.e., we can transmit long sequences of qubits reliably using $S < 1$ qubits per quantum state produced by the source.

Note that if $AB$ is prepared in a mixture of states $\rho_i \otimes |i\rangle\langle i|$ with weights $p_i$, where the $\rho_i$ are any density operators, not necessarily orthogonal, then it follows from (86), (82), and the fact that $S(|i\rangle\langle i|) = 0$ that

$$\begin{aligned} S(\sum_i p_i\rho_i \otimes |i\rangle\langle i|) &= H(p_i) + \sum_i p_i S(\rho_i \otimes |i\rangle\langle i|) \\ &= H(p_i) + \sum_i p_i S(\rho_i). \end{aligned} \tag{87}$$

The von Neumann entropy of a mixture of states $\rho_i$ with weights $p_i$, $\sum p_i\rho_i$, is a concave function of the states in the distribution, i.e.,

$$S(\sum_i p_i\rho_i) \geq \sum_i p_i S(\rho_i). \tag{88}$$

To see this, consider a composite system $AB$ in the state

$$\rho^{AB} = \sum p_i\rho_i \otimes |i\rangle\langle i|. \tag{89}$$

We have

$$S(A) = S(\sum_i p_i\rho_i) \tag{90}$$

$$S(B) = S(\sum_i p_i|i\rangle\langle i|) = H(p_i) \tag{91}$$

and

$$S(A,B) = H(p_i) + \sum_i p_i S(\rho_i) \tag{92}$$

by equation (87). By subadditivity $S(A) + S(B) \geq S(A,B)$, so:

$$S(\sum_i p_i\rho_i) \geq \sum_i p_i S(\rho_i). \tag{93}$$

It turns out that projective measurements always increase entropy, i.e., if $\rho' = \sum_i P_i\rho P_i$, then $S(\rho') \geq S(\rho)$, but generalized measurements can *decrease* entropy. Consider, for example, the generalized measurement on a qubit in the initial state $\rho$ defined by the measurement operators $M_1 = |0\rangle\langle 0|$ and $M_2 = |0\rangle\langle 1|$. (Note that these operators *do* define a generalized measurement because $M_1^\dagger M_1 + M_2^\dagger M_2 = |0\rangle\langle 0| + |1\rangle\langle 1| = I$.) After the measurement

$$\begin{aligned}\rho' &= |0\rangle\langle 0|\rho|0\rangle\langle 0| + |0\rangle\langle 1|\rho|1\rangle\langle 0| \\ &= \mathrm{Tr}(\rho)|1\rangle\langle 1| \\ &= |1\rangle\langle 1|.\end{aligned} \tag{94}$$

So $S(\rho') = 0 \leq S(\rho)$.

## 3.3 The 'No Cloning' Theorem

In § 3.1 we saw that two nonorthogonal quantum states cannot be reliably distinguished by any measurement. A *'no cloning' theorem* establishes that nonorthogonal quantum states cannot be copied. To see this, suppose there were a device $D$ that could copy any input quantum state of a system $Q$ with states in $\mathcal{H}^Q$. Suppose the initial ready state of the device $D$ is $|0\rangle \in \mathcal{H}^D$. Then we require, for any orthonormal set of input states $\{|i\rangle\}$:

$$|i\rangle|0\rangle \xrightarrow{U} |i\rangle|i\rangle \tag{95}$$

where $U$ is the unitary transformation that implements the copying process. By linearity, it then follows that for any input state $\sum_i c_i|i\rangle$:

$$(\sum_i c_i|i\rangle)|0\rangle \xrightarrow{U} \sum_i c_i|i\rangle|i\rangle \tag{96}$$

But for copying we require that:

$$(\sum_i c_i|i\rangle)|0\rangle \xrightarrow{U} (\sum_i c_i|i\rangle)(\sum_i c_i|i\rangle) \tag{97}$$

and

$$\sum_i c_i|i\rangle|i\rangle \neq (\sum_i c_i|i\rangle)(\sum_i c_i|i\rangle) = \sum_{ij} c_i c_j|i\rangle|j\rangle \tag{98}$$

unless $c_i c_j = \delta_{ij}$, which means that the device could not copy any states that are not in the orthonormal set $|i\rangle$.

Alternatively, one might note that if two states $|\psi\rangle$ and $|\phi\rangle$ could be copied, then

$$|\psi\rangle|0\rangle \quad \xrightarrow{U} \quad |\psi\rangle|\psi\rangle \tag{99}$$

$$|\phi\rangle|0\rangle \quad \xrightarrow{U} \quad |\phi\rangle|\phi\rangle \tag{100}$$

Since unitary transformations preserve inner produces, we require that

$$\langle\psi|\phi\rangle = \langle\psi|\phi\rangle\langle\psi|\phi\rangle \tag{101}$$

which is possible if and only if if$\langle\psi|\phi\rangle = 1$ or 0. That is: for cloning to be possible, either the states are identical, or they are orthogonal.

The 'no cloning' theorem was proved independently by Dieks [1982] and Wootters and Zurek [1982]. An important extension of this result to mixtures is due to Barnum, Caves, Fuchs, Jozsa, and Schumacher [1996a]. In a cloning process, a ready state $\sigma$ of a system $B$ and the state to be cloned $\rho$ of a system $A$ are transformed into two copies of $\rho$. In a more general *broadcasting* process, a ready state $\sigma$ and the state to be broadcast $\rho$ are transformed to a new state $\omega$ of $AB$, where the marginal state $\omega$ with respect to both $A$ and $B$ is $\rho$, i.e.,

$$\rho_A = \mathrm{Tr}_B(\omega) = \rho \tag{102}$$

$$\rho_B = \mathrm{Tr}_A(\omega) = \rho \tag{103}$$

The 'no cloning' theorem states that a set of pure states can be cloned if and only if the states are mutually orthogonal. The *'no broadcasting' theorem* states that an arbitrary set of states can be broadcast if and only if they are represented by mutually commuting density operators. Classically, since all pure states are, in a formal sense, orthogonal and all operators (representing real-valued functions on phase space) commute, both cloning and broadcasting are possible. Note that broadcasting reduces to cloning for pure states.

Of course, it is always possible to build a special-purpose device to clone a given (known) quantum state $|\psi\rangle$, because this would simply be a device that prepares the

state $|\psi\rangle$. The 'no cloning' theorem, from another perspective, is just the statement of the quantum measurement problem (see §7): measurements, in the classical sense of reproducing in a second system a copy of the state of the first system (or, more generally, a 'pointer state' that represents the state of the first system), are impossible in quantum mechanics, except for measurements restricted to orthogonal sets of input states.

A modification of the argument leading to Eqs. (99)–(101) shows that no information gain about the identity of nonorthogonal states is possible without disturbing the states. Suppose the device $D$ acts as a measuring device that records some information about the identity of the input state, i.e., the output state of the device is different for different input states $|\psi\rangle, |\phi\rangle$; and that the device does not disturb the input states. Then

$$|\psi\rangle|0\rangle \quad \stackrel{U}{\longrightarrow} \quad |\psi\rangle|\psi'\rangle \tag{104}$$

$$|\phi\rangle|0\rangle \quad \stackrel{U}{\longrightarrow} \quad |\phi\rangle|\phi'\rangle \tag{105}$$

from which it follows that

$$\langle\psi|\phi\rangle = \langle\psi|\phi\rangle\langle\psi'|\phi'\rangle \tag{106}$$

and so

$$\langle\psi'|\phi'\rangle = 1 \tag{107}$$

since $\langle\psi|\phi\rangle \neq 0$ if $|\psi\rangle$ and $|\phi\rangle$ are nonorthogonal. In other words, if there is no disturbance to the nonorthogonal input states, there can be no information gain about the identity of the states. So, for example, an eavesdropper, Eve, could gain no information about the identity of nonorthogonal quantum states communicated between Alice and Bob without disturbing the states, which means that passive eavesdropping is impossible for quantum information.

The observation that a set of pure states can be cloned if and only if they are mutually orthogonal is equivalent to the observation that a set of pure states can be reliably distinguished if and only if they are mutually orthogonal. For if we could distinguish a pair of states $|\psi\rangle$ and $|\phi\rangle$, then we could copy them by simply preparing the states with special-purpose preparation devices for $|\psi\rangle$ and $|\phi\rangle$. And if we could copy the states, then we could prepare as many copies as we liked of each state. Because the product states $|\psi\rangle^{\otimes n}$ and $|\phi\rangle^{\otimes n}$ become orthogonal in the limit as $n \to \infty$, these states are certainly distinguishable, and so the possibility of cloning the states $|\psi\rangle$ and $|\phi\rangle$ would provide a means of distinguishing them.

Note also that, by a similar argument, cloning would allow different mixtures associated with the same density operator to be distinguished. The equal-weight mixture of qubit states $|\uparrow_z\rangle = |0\rangle, |\downarrow_z\rangle = |1\rangle$ (the eigenstates of the spin observable $Z = \sigma_z$) has the same density operator, $I/2$, as the equal-weight mixture of states $|\uparrow_x\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle), |\downarrow_x\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ (the eigenstates of $X = \sigma_x$). Since the cloned states $|\uparrow_x\rangle^{\otimes n}, |\downarrow_x\rangle^{\otimes n}$ become distinguishable from the cloned states $|\uparrow_z\rangle^{\otimes n}$, $|\downarrow_z\rangle^{\otimes n}$, cloning would allow the two mixtures to be distinguished.

This possibility would also allow superluminal signalling. For suppose Alice and Bob shared the entangled state $\frac{1}{\sqrt{2}}(|0\rangle|1\rangle - |1\rangle|0\rangle)$. If Alice measured $X$ or $Z$ on her qubit, she would steer Bob's qubit into the mixture $\frac{1}{2}|\uparrow_x\rangle\langle\uparrow_x| + \frac{1}{2}|\downarrow_x\rangle\langle\downarrow_x|$ or the

mixture $\frac{1}{2}|\uparrow_z\rangle\langle\uparrow_z| + \frac{1}{2}|\downarrow_z\rangle\langle\downarrow_z|$. If Bob could distinguish these mixtures by cloning, in a shorter time than the time taken for light to travel between Alice and Bob, he would be able to ascertain whether Alice measured $X$ or $Z$, so 1 bit of information would be transferred from Alice to Bob superluminally.

## 3.4 Accessible Information

The ability to exploit quantum states to perform new sorts of information-processing tasks arises because quantum states have different distinguishability properties than classical states. Of course, it is not the mere lack of distinguishability of quantum states that is relevant here, but the different sort of distinguishability enjoyed by quantum states. This indistinguishability is reflected in the limited *accessibility* of quantum information.

To get a precise handle on this notion of accessibility, consider a classical information source in Shannon's sense, with Shannon entropy $H(X)$. Suppose the source produces symbols represented as the values $x$ (in an alphabet $\mathcal{X}$) of a random variable $X$, with probabilities $p_x$, and that the symbols are encoded as quantum states $\rho_x$, $x \in X$. The mutual information $H(X : Y)$ (as defined by Eqs. (22), (23), (24)) is a measure of how much information one gains, on average, about the value of the random variable $X$ on the basis of the outcome $Y$ of a measurement on a given quantum state. The *accessible information* is defined as:

$$\text{Sup } H(X\!:\!Y) \tag{108}$$

over all possible measurements.

The *Holevo bound* on mutual information provides an important upper bound to accessible information:

$$H(X\!:\!Y) \leq S(\rho) - \sum_x p_x S(\rho_x) \tag{109}$$

where $\rho = \sum_x p_x \rho_x$ and the measurement outcome $Y$ is obtained from a measurement defined by a POVM $\{E_y\}$. Since $S(\rho) - \sum_x p_x S(\rho_x) \leq H(X)$ by Eq. (86), with equality if and only if the states $\rho_x$ have orthogonal support, we have:

$$H(X\!:\!Y) \leq H(X) \tag{110}$$

Note that $X$ can be distinguished from $Y$ if and only if $H(X : Y) = H(X)$. If the states $\rho_x$ are orthogonal pure states, then in principle there exists a measurement that will distinguish the states, and for such a measurement $H(X : Y) = H(X)$. In this case, the accessible information is the same as the entropy of preparation of the quantum states, $H(X)$. But if the states are nonorthogonal, then $H(X\!:\!Y) < H(X)$ and there is no measurement, even in the generalized sense, that will enable the reliable identification of $X$.

Note, in particular, that if the values of $X$ are encoded as the pure states of a qubit, then $H(X : Y) \leq S(\rho)$ and $S(\rho) \leq 1$. It follows that at most 1 bit of information can be extracted from a qubit by measurement. If $X$ has $k$ equiprobable values,

$H(X) = \log k$. Alice could encode these $k$ values into a qubit by preparing it in an equal-weight mixture of $k$ nonorthogonal pure states, but Bob could only extract at most 1 bit of information about the value of $X$. For an $n$-state quantum system associated with an $n$-dimensional Hilbert space, $S(\rho) \leq \log n$. So even though Alice could encode any amount of information into such an $n$-state quantum system (by preparing the state as a mixture of nonorthogonal states), the most information that Bob could extract from the state by measurement is $\log n$, which is the same as the maximum amount of information that could be encoded into and extracted from an $n$-state classical system. It might seem, then, that the inaccessibility of quantum information as quantified by the Holevo bound would thwart any attempt to exploit quantum information to perform nonclassical information-processing tasks. In the following sections, we shall see that this is not the case: surprisingly, the inaccessibility of quantum information can actually be exploited in information-processing tasks that transcend the scope of classical information.

For an insightful derivation of the Holevo bound (essentially reproduced below), see [Nielsen and Chuang, 2000, Theorem 12.1, p. 531]. The basic idea is the following: Suppose Alice encodes the distinguishable symbols of a classical information source with entropy $H(X)$ as quantum states $\rho_x$ (not necessarily orthogonal). That is, Alice has a quantum system $P$, the preparation device, with an orthonormal pointer basis $|x\rangle$ corresponding to the values of the random variable $X$, which are produced by the source with probabilities $p_x$. The preparation interaction correlates the pointer states $|x\rangle$ with the states $\rho_x$ of a quantum system $Q$, so that the final state of $P$ and $Q$ after the preparation interaction is:

$$\rho^{PQ} = \sum_x p_x |x\rangle\langle x| \otimes \rho_x. \tag{111}$$

Alice sends the system $Q$ to Bob, who attempts to determine the value of the random variable $X$ by measuring the state of $Q$. The initial state of $P, Q$, and Bob's measuring instrument $M$ is:

$$\rho^{PQM} = \sum_x p_x |x\rangle\langle x| \otimes \rho_x \otimes |0\rangle\langle 0| \tag{112}$$

where $|0\rangle\langle 0|$ is the initial ready state of $M$. Bob's measurement can be described by a quantum operation $\mathcal{E}$ on the Hilbert space $\mathcal{H}^Q \otimes \mathcal{H}^M$ that stores a value of $y$, associated with a POVM $\{E_y\}$ on $\mathcal{H}^Q$, in the pointer state $|y\rangle$ of $M$, i.e., $\mathcal{E}$ is defined for any state $\sigma \in \mathcal{H}^Q$ and initial ready state $|0\rangle \in \mathcal{H}^M$ by:

$$\sigma \otimes |0\rangle\langle 0| \xrightarrow{\mathcal{E}} \sum_y \sqrt{E_y}\sigma\sqrt{E_y} \otimes |y\rangle\langle y|. \tag{113}$$

We have (recall the definition of quantum mutual information in Eqs. (78)–(80)):

$$S(P\!:\!Q) = S(P\!:\!Q, M)^8 \tag{114}$$

[8]The notation $S(P:Q,M)$ refers to the mutual information between the system $P$ and the composite system consisting of the system $Q$ and the measuring device $M$, in the initial state (111). That is, the comma notation refers to the joint system (cf. Eqs. (78), (80)): $S(P:Q,M) = S(P) - S(P|Q,M) = S(P) + S(Q,M) - S(P,Q,M)$

because $M$ is initially uncorrelated with $PQ$ and

$$S(P'\!:\!Q', M') \leq S(P\!:\!Q, M) \tag{115}$$

because it can be shown ([Nielsen and Chuang, 2000, Theorem 11.15, p. 522]) that quantum operations never increase mutual information (primes here indicate states after the application of $\mathcal{E}$). Finally:

$$S(P'\!:\!Q', M') \tag{116}$$

because discarding systems never increases mutual information ([Nielsen and Chuang, 2000, Theorem 11.15, p. 522]), and so:

$$S(P'\!:\!M') \leq S(P\!:\!Q) \tag{117}$$

which (following some algebraic manipulation) is the statement of the Holevo bound, i.e., (117) reduces to (109).

To see this, note (from (111)) that

$$\rho^{PQ} = \sum_x p_x |x\rangle\langle x| \otimes \rho_x \tag{118}$$

So $S(P) = H(p_{x)}), S(Q) = S(\sum_x p_x \rho_x) = S(\rho)$ and, by (87),

$$S(P, Q) = H(p_{x)} + \sum_x p_x S(\rho_x) \tag{119}$$

since the states $|x\rangle\langle x| \otimes \rho_x$ have support on orthogonal subspaces in $\mathcal{H}^P \otimes \mathcal{H}^Q$. It follows that

$$\begin{aligned} S(P\!:\!Q) &= S(P) + S(Q) - S(P, Q) \\ &= S(\rho) - \sum_x p_x S(\rho_x) \end{aligned} \tag{120}$$

which is the right hand side of the Holevo bound.

For the left hand side:

$$\begin{aligned} \rho^{P'M'} &= \mathrm{Tr}_{Q'}(\rho^{P'Q'M'}) && (121) \\ &= \mathrm{Tr}_{Q'}(\sum_{xy} p_x |x\rangle\langle x| \otimes \sqrt{E_y}\rho_x\sqrt{E_y} \otimes |y\rangle\langle y|) && (122) \\ &= \sum_{xy} p_x \mathrm{Tr}(E_y \rho_x E_y) |x\rangle\langle x| \otimes |y\rangle\langle y| && (123) \\ &= \sum_{xy} p(x, y) |x\rangle\langle x| \otimes |y\rangle\langle y|, && (124) \end{aligned}$$

since $p(x, y) = p_x p(y \mid x) = p_x \mathrm{Tr}(\rho_x E_y) = p_x \mathrm{Tr}(\sqrt{E_y}\rho_x\sqrt{E_y})$, and so $S(P' : M') = H(X : Y)$.

The Holevo bound limits the representation of classical bits by qubits. Putting it another way, the Holevo bound characterizes the resource cost of encoding classical bits as qubits: one qubit is necessary and sufficient. Can we represent qubits by bits? If so, what is the cost of a qubit in terms of bits? This question is answered by the following result [Barnum *et al.*, 2001]: A quantum source of nonorthogonal signal states can be compressed with arbitarily high fidelity to $\alpha$ qubits per signal plus any number of classical bits per signal if and only if $\alpha$ is at least as large as the von Neumann entropy $S$ of the source. This means that a generic quantum source cannot be separated into a classical and quantum part: quantum information cannot be traded for any amount of classical information.

## 3.5 Quantum Information Compression

As pointed out in §3.2, Shannon's source coding theorem (noiseless channel coding theorem) and the core notion of a typical sequence can be generalized for quantum sources. This was first shown by Jozsa and Schumacher [1994] and Schumacher [1995]. See also [Barnum *et al.*, 1996b].

For a classical information bit source, where the output of the source is given by a random variable $X$ with two possible values $x_1, x_2$ with probabilities $p_1, p_2$, the Shannon entropy of the information produced by the source is $H(X) = H(p_1, p_2)$. So by Shannon's source coding theorem the information can be compressed and communicated to a receiver with arbitrarily low probability of error by using $H(X)$ bits per signal, which is less than one bit if $p_1 \neq p_2$.

Now suppose the source produces qubit states $|\psi_1\rangle, |\psi_2\rangle$ with probabilities $p_1, p_2$. The Shannon entropy of the mixture $\rho = p_1|\psi_1\rangle\langle\psi_1| + p_2|\psi_2\rangle\langle\psi_2|$ is $S(\rho)$. Schumacher's generalization of Shannon's source coding theorem shows that the quantum information encoded in the mixture $\rho$ can be compressed and communicated to a receiver with arbitrarily low probability of error by using $S(\rho)$ qubits per signal, and $S(\rho) < 1$ if the qubit states are nonorthogonal.

Note that the signals considered here are *qubit states*. What Schumacher's theorem shows is that we can reliably communicate the sequence of qubit states produced by the source by sending less than one qubit per signal. Note also that since $S(\rho) < H(p_1, p_2)$ if the qubit states are nonorthogonal, the quantum information represented by the sequence of qubit states can be compressed beyond the classical limit of the classical information associated with the entropy of preparation of $\rho$ (i.e., the Shannon entropy of the random variable whose values are the labels of the qubit states).

Since the individual states in a mixture are not in general distinguishable, there are two distinct sorts of compression applicable to quantum information that do not apply to classical information. In *blind compression*, the sequence of quantum states produced by a source is compressed via a compression scheme that depends only on the identities of the quantum states and their probabilities, i.e., the input to the compression scheme is the density operator associated with the distribution. In *visible compression*, the identity of each individual quantum state produced by the source is assumed to be known, i.e., the input to the compression scheme is an individual quantum state in the sequence produced by the source, and the compression of the state is based on the probability distribution of such states.

An example: an inefficient visible compression scheme of the above qubit source ($|\psi_1\rangle, |\psi_2\rangle$ with probabilities $p_1, p_2$) would simply involve sending the classical information of the quantum state labels, compressed to $H(p_1, p_2)$ bits per signal, to the receiver, where the original qubit states could then be prepared after decompression of the classical information. This scheme is not optimal by Schumacher's theorem (for nonorthogonal qubit states) because $S(\rho) < H(p_1, p_2)$. Of course, Schumacher's theorem refers to a compression rate of $S(\rho)$ *qubits* per quantum signal, while the application of Shannon's theorem here refers to $H(p_2, p_2)$ *bits* per classical signal. But note that the communication of one classical bit requires the same physical resource as the communication of one qubit, prepared in one of two orthogonal basis states. Note also that sending the (nonorthogonal) qubit states themselves, which would require one qubit per signal, would not convey the identity of the states in the sequence to the transmitter. So the classical information about the individual state labels in the sequence (which would be bounded by $\log n$ per signal if we considered a source producing $n$ qubit states) is really redundant if the aim is to communicate the quantum information associated with the sequence of qubit states.

Remarkably, Schumacher's theorem shows that the optimal compressibility of the quantum information associated with a sequence of quantum pure states is $S(\rho)$ qubits per signal, for blind *or* visible compression.

To see the general idea, consider a source of (possibly nonorthogonal) qubits $|\psi_1\rangle, |\psi_2\rangle$ with probabilities $p_1, p_2$. The density operator of the probability distribution is $\rho = p_1|\psi_1\rangle\langle\psi_1| + p_2|\psi_2\rangle\langle\psi_2|$.

An $n$-sequence of states produced by the source is represented by a state

$$|\Psi_{i_1 \ldots i_n}\rangle = |\psi_{i_1}\rangle \ldots |\psi_{i_n}\rangle \tag{125}$$

in $\mathcal{H}_2^{\otimes n}$. Each such state has a probability $p_{i_1 \ldots i_n} = p_{i_1} \ldots p_{i_n}$. The $n$-sequences span the $2^n$-dimensional Hilbert space $\mathcal{H}_2^{\otimes n}$, but as $n \to \infty$ it turns out that the probability of finding an $n$-sequence in a 'typical subspace' (in a measurement, on an $n$-sequence produced by the source, of the projection operator onto the subspace ) tends to 1. That is, for any $\epsilon, \delta > 0$, there is a subspace $\mathcal{T}_\delta^{(n)}$ of dimension between $2^{n(S(\rho)-\delta)}$ and $2^{n(S(\rho)+\delta)}$, with projection operator $P_\delta^{(n)}$, such that:

$$\sum_{\text{all sequences}} p_{i_1 \ldots i_n} \mathrm{Tr}(|\Psi_{i_1 \ldots i_n}\rangle\langle\Psi_{i_1 \ldots i_n}|P_\delta^{(n)}) = \mathrm{Tr}(\rho^{\otimes n} P_\delta^{(n)}) \geq 1 - \epsilon. \tag{126}$$

Here $\rho^{\otimes n} = \rho \otimes \rho \ldots \rho$, the $n$-fold tensor product of $\rho$ with itself, is the density operator of $n$-sequences of states produced by the source:

$$\rho^{\otimes n} = \sum_{\text{all } n\text{-sequences}} p_{i_1 \ldots i_n} |\Psi_{i_1 \ldots i_n}\rangle\langle\Psi_{i_1 \ldots i_n}| \tag{127}$$

$$= \sum_{\text{all } n\text{-sequences}} p_{i_1} \ldots p_{i_n} |\psi_{i_1}\rangle\langle\psi_{i_1}| \otimes \ldots \otimes |\psi_{i_n}\rangle\langle\psi_{i_n}| \tag{128}$$

where each state $|\psi_{i_j}\rangle$ is one of $k$ possible states in a $d$-dimensional Hilbert space. Recall that the statistical properties of such $n$-sequences of states, for all possible measurements, is given by $\rho^{\otimes n}$ and does not depend on the representation of $\rho^{\otimes n}$ as a

particular mixture of states. Since $S(\rho) \leq 1$ for a qubit source, the dimension of $\mathcal{T}_\delta^{(n)}$ decreases exponentially in $\mathcal{H}_2^{\otimes n}$ as $n \to \infty$, i.e., the typical subspace is exponentially small in $\mathcal{H}_2^{\otimes n}$ for large $n$.

Note that this does *not* mean that almost all $n$-sequences of states produced by the source lie in the typical subspace. Rather, almost all $n$-sequences produced by the source are such that a measurement of $P_\delta^{(n)}$ on the sequence will yield the value 1, i.e., almost all $n$-sequences produced by the source will answer 'yes' in a measurement of the projection operator onto the typical subspace. So, in this sense, most sequences produced by the source will be found to lie in the typical subspace on measurement, and for any subspace $\mathcal{V}$ of dimension less than $2^{n(S(\rho)-\delta)}$ it can be shown that the average probability of finding an $n$-sequence produced by the source in $\mathcal{V}$ is less than any pre-assigned $\epsilon$ for sufficiently large $n$.

Consider now the general case where the source produces $k$ states $|\psi_1\rangle \ldots, |\psi_k\rangle \in \mathcal{H}_d$ (not necessarily orthogonal) with probabilities $p_1 \ldots, p_k$. Here the density operator associated with the source is $\rho = \sum_{i=1}^k p_i |\psi_i\rangle\langle\psi_i|$. Sequences of length $n$ span a subspace of $d^n = 2^{n \log d}$ dimensions and the typical subspace $\mathcal{T}_\delta^{(n)}$ has dimension between $2^{n(s(\rho)-\delta)}$ and $2^{n(s(\rho)+\delta)}$, which is again exponentially small in $\mathcal{H}_d^{\otimes n}$ because $S(\rho) \leq \log d$.

For comparison with Shannon's theorem, we write $\rho$ in the spectral representation as:

$$\rho = \sum_x p(x)|x\rangle\langle x| \tag{129}$$

where $\{p(x)\}$ is the set of non-zero eigenvalues of $\rho$ and $\{|x\rangle\}$ is an orthonormal set of eigenstates of $\rho$. If $\rho$ has eigenvalues $p(x)$ and eigenstates $|x\rangle$, then $\rho^{\otimes n}$ has eigenvalues $p(x_1)p(x_2)\ldots p(x_n)$ and eigenstates $|x_1\rangle|x_2\rangle\ldots|x_n\rangle$.

A $\delta$-typical state is defined as a state $|x_1\rangle|x_2\rangle\ldots|x_n\rangle$ for which the sequence $x_1, x_2, \ldots, x_n$ is a $\delta$-typical sequence, in the sense that (cf. Eq. (13)):

$$2^{-n(S(\rho)+\delta)} < p(x_1 \ldots x_n) < 2^{-n(S(\rho)-\delta)}. \tag{130}$$

The $\delta$-typical subspace $\mathcal{T}_\delta^{(n)}$ is the subspace spanned by all the $\delta$-typical states. Denote the projection operator onto $\mathcal{T}_\delta^{(n)}$ by:

$$P_\delta^{(n)} = \sum_{\delta-\text{typical states}} |x_1\rangle\langle x_1| \otimes |x_2\rangle\langle x_2| \ldots |x_n\rangle\langle x_n| \tag{131}$$

Then, for a fixed $\delta > 0$, it can be shown that for any $\epsilon > 0$ and sufficiently large $n$

$$\mathrm{Tr}(P_\delta^{(n)}\rho^{\otimes n}) \geq 1 - \epsilon; \tag{132}$$

and the dimension of $\mathcal{T}_\delta^{(n)}$ $(= \mathrm{Tr}(P_\delta^{(n)}))$ satisfies

$$(1-\epsilon)2^{n(S(\rho)-\delta)} \leq \dim \mathcal{T}_\delta^{(n)} \leq 2^{n(S(\rho)+\delta)}. \tag{133}$$

That is, the dimension of $\mathcal{T}_\delta^{(n)}$ is roughly $2^{nS(\rho)}$, which is exponentially smaller than the dimension of $\mathcal{H}^{\otimes n}$, as $n \to \infty$.

It follows that the density operator $\rho^{\otimes n}$ can be replaced with a density operator $\tilde{\rho}^{\otimes n}$ with support on the typical subspace (take $\rho^{\otimes n}$ in the spectral representation, where the matrix is diagonal with $2^{n \log d}$ eigenvalues $p(x_1 \ldots x_n) = p(x_1) \ldots p(x_n)$, and replace all $p(x_1 \ldots x_n)$ that do not correspond to typical sequences with zeros).

Before considering a compression/decompression scheme for quantum information, we need a measure of the reliability of such a scheme in terms of the fidelity, as in the case of classical information. The following definition generalizes the classical notion of fidelity in §2 (see [Jozsa, 1998, p. 70]): If $|\psi\rangle$ is any pure quantum state and $\rho$ any mixed state, the fidelity between $\rho$ and $|\psi\rangle$ is:

$$F(\rho, |\psi\rangle) = \mathrm{Tr}((\rho|\psi\rangle\langle\psi|)) = \langle\psi|\rho|\psi\rangle \tag{134}$$

which is the probability that a measurement of the projection operator $|\psi\rangle\langle\psi|$ in the state $\rho$ yields the outcome 1, i.e., it is the probability that $\rho$ passes a test of being found to be $|\psi\rangle$ on measurement. Note that for a pure state $\rho = |\psi\rangle\langle\psi|$, $F(|\phi\rangle, |\psi\rangle) = |\langle\psi|\phi\rangle|^2$. The fidelity between two mixed states $\rho$ and $\sigma$ is defined as:[9]

$$F(\rho, \sigma) = \max|\langle\psi|\phi\rangle|^2 = (\mathrm{Tr}(\sqrt{\rho^{1/2}\sigma\rho^{1/2}}))^2. \tag{135}$$

for all purifications $|\psi\rangle$ of $\rho$ and $|\phi\rangle$ of $\sigma$. Note that in spite of appearances, $F(\rho, \sigma)$ is symmetric in $\rho$ and $\sigma$.

In the case of a source of $n$-sequences of quantum states $|\Psi_{i_1 \ldots i_n}\rangle = |\psi_{i_1}\rangle \ldots |\psi_{i_n}\rangle$ with prior probabilities $p_{i_1 \ldots i_n} = p_{i_1} \ldots p_{i_n}$, a compression/decompression scheme will in general yield a mixed state $\rho_{i_1 \ldots i_n}$. The average fidelity of a compression-decompression scheme for an $n$-sequence of quantum states is defined as:

$$F_n = \sum_{\text{all } n\text{-sequences}} p_{i_1 \ldots i_n} \mathrm{Tr}(\rho_{i_1 \ldots i_n} |\Psi_{i_1 \ldots i_n}\rangle\langle\Psi_{i_1 \ldots i_n}|) \tag{136}$$

Schumacher's quantum source coding theorem (or quantum noiseless channel coding theorem) for a quantum source that produces quantum states $|\psi_1\rangle \ldots |\psi_n\rangle \in \mathcal{H}_d$ with probabilities $p_1 \ldots p_n$ (so the density operator corresponding to the output of the source is $\rho = \sum p_i |\psi_i\rangle\langle\psi_i|$), states that

> for any $\epsilon, \delta > 0$: (i) there exists a compression/decompression scheme using $S(\rho) + \delta$ qubits per state for $n$-length sequences produced by the source that can be decompressed by the receiver with fidelity $F_n > 1 - \epsilon$, for sufficiently large $n$, and (ii) any compression/decompression scheme using $S(\rho) - \delta$ qubits per state for $n$-length sequences will have a fidelity $F_n < \epsilon$, for sufficiently large $n$.

A compression/decompression scheme for such a quantum source would go as follows: The transmitter applies a unitary transformation $U$ in $\mathcal{H}_d^{\otimes n}$ (dimension =

[9] Note that Nielsen and Chuang [2000, p. 409] define the fidelity $F(\rho, \sigma)$ as the square root of the quantity defined here. If $\rho$ and $\sigma$ commute, they can be diagonalized in the same basis. The definition then reduces to their definition of the classical fidelity between two probablity distributions defined by the eigenvalues of $\rho$ and $\sigma$ in footnote 3 in §2.1.

$d^n = 2^{n \log d}$) which maps any state in the typical subspace onto a linear superposition of sequences of $n \log d$ qubits, where all but the first $nS(\rho)$ qubits are in the state $|0\rangle$, and then transmits the first $nS(\rho)$ qubits to the receiver. So the transmitter compresses $n \log d$ qubits to $nS(\rho)$ qubits. The receiver adds $n \log d - nS(\rho)$ qubits in the state $|0\rangle$ and applies the unitary transformation $U^{-1}$. Since the initial $nS(\rho)$ qubits will in general be slightly entangled with the remaining $n \log d - nS(\rho)$ qubits, discarding these qubits amounts to tracing over the associated dimensions, so replacing these qubits with the state $|0\rangle$ will produce a mixed state $\tilde{\rho_n}$. The state $U^{-1}\tilde{\rho_n}$ will pass a test of being found to be the original state $|\Psi_{i_1 \ldots i_n}\rangle$ with fidelity greater than $1 - \epsilon$.

# 4 Entanglement Assisted Quantum Communication

In this section I show how entanglement can be exploited as a channel for the reliable transmission of quantum information. I discuss two related forms of entanglement assisted communication: quantum teleportation in §4.1 and quantum dense coding in §4.2.

## 4.1 Quantum Teleportation

As mentioned in §1, Schrödinger introduced the term 'entanglement' to describe the peculiar nonlocal correlations of the EPR-state in an extended two-part commentary [1935, 1936] on the Einstein-Podolsky-Rosen argument [Einstein *et al.*, 1935]. Schrödinger regarded entangled states as problematic because they allow the possibility of what he called 'remote steering,' which he regarded as a mathematical artefact of the Hilbert space theory and discounted as a physical possibility. As it turns out, quantum teleportation is an experimentally confirmed application of remote steering between two separated systems. This was first pointed out in a paper by Bennett, Brassard, Crépeau, Jozsa, Peres, and Wootters [1993] and later experimentally confirmed by several groups using a variety of different techniques [Bouwmeester *et al.*, 1997, Boschi *et al.*, 1998, Furasawa *et al.*, 1998, Nielsen *et al.*, 1998].

In the 1935 paper, Schrödinger considered pure entangled states with a unique biorthogonal decomposition, as well as cases like the EPR-state, where a biorthogonal decomposition is non-unique. He showed that suitable measurements on one system can fix the (pure) state of the entangled distant system, and that this state depends on what observable one chooses to measure, not merely on the outcome of that measurement. In the second paper, he showed that a 'sophisticated experimenter,' by performing a suitable local measurement on one system, can 'steer' the distant system into any mixture of pure states represented by its reduced density operator. So the distant system can be steered (probabilistically, depending on the outcome of the local measurement) into any pure state in the support of the reduced density operator, with a nonzero probability that depends only on the pure state. For a mixture of linearly independent states of the distant system, the steering can be done by performing a local standard projection-valued measurement in a suitable basis. If the states are linearly dependent, the experimenter performs a generalized measurement (associated with a POVM), which amounts to enlarging the experimenter's Hilbert space by adding an

ancilla, so that the dimension of the enlarged Hilbert space is equal to the number of linearly independent states. As indicated in §3.1.1, Schrödinger's analysis anticipated the later result by Hughston, Jozsa, and Wootters [1993].

Suppose Alice and Bob, the traditional protagonists in any two-party communication protocol, each holds one of a pair of qubits in the entangled state:

$$|\Psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle_A|1\rangle_B - |1\rangle_A|0\rangle_B) \tag{137}$$

Bob's qubit separately is in the mixed state $\rho_B = I/2$, which can be interpreted as an equal weight mixture of the orthogonal states $|0\rangle_B$, $|1\rangle_B$, or, equivalently, as an infinity of other mixtures including, to take a specific example, the equal weight mixture of the four nonorthogonal normalized states:

$$\begin{aligned}
|\phi_1\rangle_B &= \alpha|0\rangle_B + \beta|1\rangle_B \\
|\phi_2\rangle_B &= \alpha|0\rangle_B - \beta|1\rangle_B \\
|\phi_3\rangle_B &= \beta|0\rangle_B + \alpha|1\rangle_B \\
|\phi_4\rangle_B &= \beta|0\rangle_B - \alpha|1\rangle_B
\end{aligned}$$

That is:

$$\rho_B = I/2 = \frac{1}{4}(|\phi_1\rangle\langle\phi_1| + |\phi_2\rangle\langle\phi_2| + |\phi_3\rangle\langle\phi_3| + |\phi_4\rangle\langle\phi_4|) \tag{138}$$

If Alice measures the observable with eigenstates $|0\rangle_A$, $|1\rangle_A$ on her qubit A, and Bob measures the corresponding observable on his qubit B, Alice's outcomes will be oppositely correlated with Bob's outcomes (0 with 1, and 1 with 0). If, instead, Alice prepares an ancilla qubit $A'$ in the state $|\phi_1\rangle_{A'} = \alpha|0\rangle_{A'} + \beta|1\rangle_{A'}$ and measures an observable on the pair of qubits $A' + A$ in her possession with eigenstates:

$$|1\rangle = (|0\rangle_{A'}|1\rangle_A - |1\rangle_{A'}|0\rangle_A)/\sqrt{2} \tag{139}$$
$$|2\rangle = (|0\rangle_{A'}|1\rangle_A + |1\rangle_{A'}|0\rangle_A)/\sqrt{2} \tag{140}$$
$$|3\rangle = (|0\rangle_{A'}|0\rangle_A - |1\rangle_{A'}|1\rangle_A)/\sqrt{2} \tag{141}$$
$$|4\rangle = (|0\rangle_{A'}|0\rangle_A + |1\rangle_{A'}|1\rangle_A)/\sqrt{2} \tag{142}$$

(the Bell states defining the Bell basis in $\mathcal{H}^{A'} \otimes \mathcal{H}^A$), she will obtain the outcomes 1, 2, 3, 4 with equal probability of 1/4, and these outcomes will be correlated with Bob's states $|\phi_1\rangle_B$, $|\phi_2\rangle_B$, $|\phi_3\rangle_B$, $|\phi_4\rangle_B$. That is, if Bob checks to see whether his particle is in the state $|\phi_i\rangle_B$ when Alice reports that she obtained the outcome $i$, he will find that this is always in fact the case. This follows because

$$|\phi_1\rangle_{A'}|\Psi\rangle = \frac{1}{2}(-|1\rangle|\phi_1\rangle_B - |2\rangle|\phi_2\rangle_B + |3\rangle|\phi_3\rangle_B - |4\rangle|\phi_4\rangle_B) \tag{143}$$

In this sense, Alice can steer Bob's particle into any mixture compatible with his density operator $\rho_B = I/2$ by an appropriate local measurement.

What Schrödinger found problematic about entanglement was the possibility of remote steering in the above sense [1935, p. 556]:

> It is rather discomforting that the theory should allow a system to be steered or piloted into one or the other type of state at the experimenter's mercy in spite of his having no access to it.

Now, remote steering in this probabilistic sense is precisely what makes quantum teleportation possible. Suppose Alice and Bob share a pair of qubits in the entangled state (137) and Alice is given a qubit $A'$ in an *unknown* state $|\phi_1\rangle$ that she would like to send to Bob. There is no procedure by which Alice can determine the identity of the unknown state, but even if she could, the amount of classical information that Alice would have to send to Bob in order for him to prepare the state $|\phi_1\rangle$ is potentially infinite, since the precise specification of a general normalized qubit state $\alpha|0\rangle + \beta|1\rangle$ requires two real parameters ( the number of independent parameters is reduced from four to two because $|\alpha|^2 + |\beta|^2 = 1$ and the overall phase is irrelevant). Alice could send the qubit itself to Bob, but the quantum information in the qubit state might be corrupted by transmission through a possibly noisy environment.

Instead, for the cost of just two bits of classical information, Alice can succeed in communicating the unknown quantum state $|\phi_1\rangle$ to Bob with perfect reliability. What Alice does is to measure the 2-qubit system $A' + A$ in her possession in the Bell basis. Depending on the outcome of her measurement, $i = 1, 2, 3$, or 4 with equal probability, Bob's qubit will be steered into one of the states $|\phi_1\rangle_B$, $|\phi_2\rangle_B$, $|\phi_3\rangle_B$, $|\phi_4\rangle_B$. If Alice communicates the outcome of her measurement to Bob (requiring the transmission of two bits of classical information), Bob can apply one of four local unitary transformations in his Hilbert space to obtain the state $|\phi_1\rangle_B$:

$i = 1$: do nothing, i.e., apply the identity transformation $I$

$i = 2$: apply the transformation $\sigma_z$

$i = 3$: apply the transformation $\sigma_x$

$i = 4$: apply the transformation $i\sigma_y$

where $\sigma_x, \sigma_y, \sigma_z$ are the Pauli spin matrices.

The trick that results in the communication of the state $|\phi_1\rangle$ from Alice to Bob, without the qubit $A'$ literally traveling from Alice to Bob, is the ability afforded Alice by the shared entangled state to correlate one of four measurement outcomes (each occurring with probability 1/4) with one of four states that together represent a particular decomposition of Bob's mixed state. The communication of the state of $A'$ is completed by Bob's operation, which requires that Alice sends the two bits of classical information about her measurement outcome to Bob. In the teleportation protocol, the state of the particle $A'$ is destroyed by Alice's measurement and re-created as the state of Bob's particle by Bob's operation—in fact, the systems $A$ and $A'$ end up in an entangled state as the result of Alice's measurement. Note that if the state $|\phi_1\rangle$ of $A'$ were not destroyed there would be two copies of the state, which would violate the quantum 'no cloning' theorem. So neither Alice nor Bob, nor any other party, can gain any information about the identity of the teleported state, because the recording of such information in the state of another quantum system would amount to a partial copying of the information in the teleported state.

Shared entanglement provides a secure and reliable channel for quantum communication. This might be useful for the communication of quantum information between parties in a cryptographic protocol, or for the transmission of quantum information between the processing components of a quantum computer. It is a feature of an entangled state shared by two parties that the entanglement is not affected by noise in the environment between them. So the reliability of the communication of quantum information by teleportation depends on the reliability of the required classical communication, which can be protected against noise by well-known techniques of error-correcting codes. An entangled state shared by two parties is also unaffected by changes in their relative spatial location. So Alice could teleport a quantum state to Bob without even knowing Bob's location, by broadcasting the two bits of information.

## 4.2 Quantum Dense Coding

We know from the Holevo bound (see §3.4) that the maximum amount of classical information that can be reliably communicated by encoding the information in the quantum state of a qubit is one bit, even though an arbitrarily large amount of classical information can be encoded in the state of a qubit (by encoding symbols as nonorthogonal quantum states). Quantum dense coding is a procedure, first pointed out by Bennett and Wiesner [1992], for exploiting entanglement to double the amount of classical information that can be communicated by a qubit.

Consider again the Bell states:

$$|1\rangle = (|0\rangle|1\rangle - |1\rangle|0\rangle)/\sqrt{2} \tag{144}$$

$$|2\rangle = (|0\rangle|1\rangle + |1\rangle|0\rangle)/\sqrt{2} \tag{145}$$

$$|3\rangle = (|0\rangle|0\rangle - |1\rangle|1\rangle)/\sqrt{2} \tag{146}$$

$$|4\rangle = (|0\rangle|0\rangle + |1\rangle|1\rangle)/\sqrt{2} \tag{147}$$

Suppose Alice and Bob share a pair of qubits in the state

$$|1\rangle = (|0\rangle_A|1\rangle_B - |1\rangle_A|0\rangle_B)/\sqrt{2} \tag{148}$$

By performing one of four *local* operations on the qubit in her possession defined by the unitary transformations in $\mathcal{H}^A$:

$$U_1 = I \tag{149}$$

$$U_2 = \sigma_x \tag{150}$$

$$U_3 = \sigma_x \tag{151}$$

$$U_4 = i\sigma_y \tag{152}$$

Alice can transform the state $|1\rangle$ of the qubit pair into any Bell state. For example:

$$I|1\rangle = |1\rangle \tag{153}$$

$$\sigma_z|1\rangle = |2\rangle \tag{154}$$

$$\sigma_x|1\rangle = |3\rangle \tag{155}$$

$$i\sigma_y|1\rangle = |4\rangle \tag{156}$$

So to communicate two classical bits to Bob, Alice applies one of the four operations above to her qubit and sends the qubit to Bob. Bob then performs a measurement on the two qubits in the Bell basis. Since these are orthogonal states, he can distinguish the states and identify Alice's operation.

# 5 Quantum Cryptography

Over the past few years, quantum cryptography as emerged as perhaps the most successful area of application of quantum information theoretic ideas. The main results have been a variety of provably secure protocols for key distribution, following an original proposal by Bennett and Brassard [1984], and an important 'no go' theorem by Mayers [1996b, 1997] and Lo and Chau [1998]: the impossibility of unconditionally secure two-party quantum bit commitment. The quantum bit commitment theorem generalizes previous results restricted to one-way communication protocols by Mayers [1996a] and by Lo and Chau [1997] and applies to quantum, classical, and quantum-classical hybrid schemes (since classical information, as we have seen, can be regarded as quantum information subject to certain constraints). The restriction to two-party schemes excludes schemes that involve a trusted third-party or trusted channel properties, and the restriction to schemes based solely on the principles of quantum mechanics excludes schemes that exploit special relativistic signalling constraints, or schemes that might involve time machines or the thermodynamics of black holes, etc.

In §5.1, I show how the security of quantum key distribution depends on features of quantum information—no cloning, no information gain without disturbance, entanglement—that prevent an eavesdropper from secretly gaining information about the quantum communication between two parties, i.e., completely undetectable eavesdropping is in principle impossible for quantum communication. In §5.2, I dscuss quantum bit commitment and show why unconditionally secure quantum bit commitment is impossible.

## 5.1 Key Distribution

### 5.1.1 Quantum Key Distribution Protocols

In a quantum key distribution protocol, the object is for two parties, Alice and Bob, who initially share no information, to exchange information via quantum and classical channels, so as to end up sharing a secret key which they can then use for encryption, in such a way as to ensure that any attempt by an eavesdropper, Eve, to gain information about the secret key will be detected with non-zero probability.

The one-time pad provides a perfectly secure way for Alice and Bob to communicate classical information, but this is also the *only* way that two parties can achieve perfectly security classical communication. The one-time pad is, essentially, a random sequence of bits. If Alice and Bob both have a copy of the one-time pad, Alice can communicate a message to Bob securely by converting the message to an $n$-bit binary number (according to some scheme known to both Alice and Bob), and adding (bitwise, modulo 2) the sequence of bits in the binary number to an $n$-length sequence

of bits from the top of the one-time pad. Alice sends the encrypted sequence to Bob, which Bob can then decrypt using the same sequence of bits from his copy of the one-time pad. Since the encrypted message is random, it is impossible for Eve to decrypt the message without a copy of the one-time pad. It is essential to the security of the scheme that the $n$ random bits used to encrypt the message are discarded once the message is transmitted and decrypted, and that a unique random sequence is used for each distinct message—hence the term 'one-time pad.'

This procedure guarantees perfect privacy, so long as Alice and Bob, and only Alice and Bob, can each be assumed to possess a copy of an arbitrarily long one-time pad. But this means that in order for two parties to communicate secretly, they must already share a secret: the random key. The *key distribution problem* is the problem of how to distribute the key securely in the first place without the key being secretly intercepted during transmission and copied, and the *key storage problem* is the problem of how to store the key securely without it being secretly copied. We would like a procedure that can be guaranteed to be secure against passive eavesdropping, so that Alice and Bob can be confident that their communications are in fact private.

The key idea in *quantum* cryptography is to exploit the indistinguishability of nonorthogonal quantum states, which we saw in §3.3 entails that any information gained by Eve about the identity of such states will introduce some disturbance of the states that can be detected by Alice and Bob, and the 'no cloning' theorem, which makes it impossible for Eve to copy quantum communications between Alice and Bob and store them for later analysis (perhaps using, in addition, intercepted classical communications between Alice and Bob).

A large variety of quantum key distribution schemes have been proposed following the original Bennett and Brassard protocol [1984], now known as BB84. The core idea there was for Alice to send Bob a sequence of qubits, prepared with equal probability in one of the states $|0\rangle, |1\rangle, |+\rangle, |-\rangle$, where the pair of orthogonal states $|0\rangle, |1\rangle$ are nonorthogonal to the pair of orthogonal states $|+\rangle, |-\rangle$. Bob measures each qubit randomly in either the basis $|0\rangle, |1\rangle$ or the basis $|+\rangle, |-\rangle$. Following his measurements, he publicly broadcasts the basis he used for each qubit in the sequence, and Alice publicly broadcasts which of these bases is the same as the basis she used to prepare the qubit. Alice and Bob then discard the qubits for which their bases disagree. Since the outcome states of Bob's measurements are the same as the states Alice prepared, Alice and Bob share a random key on the remaining qubits. They can then sacrifice a portion of these qubits to detect eavesdropping. Alice publicly announces the qubit state she prepared and Bob checks his measurement outcome to confirm this. If they agree on a sufficient number of qubit states (depending on the expected error rate), they conclude that there has been no eavesdropping and use the remaining portion as the secret key. If they don't agree, they conclude that the qubits have been disturbed by eavesdropping, in which case they discard all the qubits and begin the procedure again. The actual protocol involves further subtleties in which a perfectly secure secret key is distilled from the 'raw key' obtained in this way by techniques of error correction and privacy amplification.

The BB84 scheme solves the key distribution problem, in the sense that Alice and Bob, who initially share no secrets, can end up sharing a secret key via a key distribution protocol that excludes the possibility of eavesdropping, with arbitrarily high

reliability(since the length of the sequence of qubits sacrificed to detect eavesdropping can be arbitrarily long). Clearly, it does not solve the key storage problem, since the output of the key distribution protocol is stored as classical information, which is subject to passive eavesdropping.

A scheme proposed by Ekert [1991] allows Alice and Bob to create a shared random key by performing measurements on two entangled qubits. Suppose Alice and Bob share many copies of an entangled pure state of two qubits, say the Bell state $\frac{1}{\sqrt{2}}(|0\rangle|1\rangle - |1\rangle|0\rangle)$ (perhaps emitted by a common source of entangled pairs between Alice and Bob). Alice and Bob agree on three observables that they each measure on their qubits, where the measurements are chosen randomly and independently for each qubit. After a sequence of measurements on an appropriate number of pairs, Alice and Bob announce the directions of their measurements publicly and divide the measurements into two groups: those in which they measured the spin in different directions, and those in which they measured the spin in the same direction. They publicly reveal the outcomes of the first group of measurements and use these to check that the singlet states have not been disturbed by eavesdropping. Essentially, they calculate a correlation coefficient: any attempt by an eavesdropper, Eve, to monitor the particles will disturb the entangled state and result in a correlation coefficient that is bounded by Bell's inequality and is therefore distinguishable from the correlation coefficient for the entangled state. If Alice and Bob are satisfied that no eavesdropping has occurred, they use the second group of oppositely correlated measurement outcomes as the key.

### 5.1.2 Quantum Key Distribution via Pre- and Post-Selection

The Ekert scheme solves the key distribution problem as well as the key storage problem, because a new key is generated for each message from the stored entangled states, and there is no information about the key in the entangled states. Here I describe a key distribution protocol that also involves entangled states (see [Bub, 2001b]), but with a different type of test for eavesdropping. Instead of a statistical test based on Bell's theorem, the test exploits conditional statements about measurement outcomes generated by pre- and post-selected quantum states.

The peculiar features of pre- and post-selected quantum states were first pointed out by Aharonov, Bergmann, and Lebowitz [1964]. If:

(i) Alice prepares a system in a certain state $|\text{pre}\rangle$ at time $t_1$,

(ii) Bob measures some observable $M$ on the system at time $t_2$,

(iii) Alice measures an observable of which $|\text{post}\rangle$ is an eigenstate at time $t_3$, and post-selects for $|\text{post}\rangle$,

then Alice can assign probabilities to the outcomes of Bob's $M$-measurement at $t_2$, conditional on the states $|\text{pre}\rangle$ and $|\text{post}\rangle$ at times $t_1$ and $t_3$, respectively, as follows [Aharonov *et al.*, 1964, Vaidman *et al.*, 1987]:

$$\text{prob}(q_k) = \frac{|\langle \text{pre}|P_k|\text{post}\rangle|^2}{\sum_i |\langle \text{pre}|P_i|\text{post}\rangle|^2} \tag{157}$$

where $P_i$ is the projection operator onto the $i$'th eigenspace of $M$. Notice that (157)—referred to as the 'ABL-rule' (Aharonov-Bergmann-Lebowitz rule) in the following—is time-symmetric, in the sense that the states $|\text{pre}\rangle$ and $|\text{post}\rangle$ can be interchanged.

If $M$ is unknown to Alice, she can use the ABL-rule to assign probabilities to the outcomes of various hypothetical $M$-measurements. The interesting peculiarity of the ABL-rule, by contrast with the usual Born rule for pre-selected states, is that it is possible—for an appropriate choice of observables $M$, $M'$, ..., and states $|\text{pre}\rangle$ and $|\text{post}\rangle$—to assign unit probability to the outcomes of a set of mutually *noncommuting* observables. That is, Alice can be in a position to assert a conjunction of conditional statements of the form: 'If Bob measured $M$, then the outcome must have been $m_i$, with certainty, and if Bob measured $M'$, then the outcome must have been $m'_j$, with certainty, ...,' where $M, M', \ldots$ are mutually noncommuting observables. Since Bob could only have measured at most one of these noncommuting observables, Alice's conditional information does not, of course, contradict quantum mechanics: she only knows the eigenvalue $m_i$ of an observable $M$ if she knows that Bob in fact measured $M$.

Vaidman, Aharonov, and Albert [1987] discuss a case of this sort, where the outcome of a measurement of any of the three spin observables $X = \sigma_x, Y = \sigma_y, Z = \sigma_z$ of a spin-$\frac{1}{2}$ particle can be inferred from an appropriate pre- and post-selection. Alice prepares the Bell state

$$|\text{pre}\rangle = \frac{1}{\sqrt{2}}(|\uparrow_z\rangle_A|\uparrow_z\rangle_C + |\downarrow_z\rangle_A|\downarrow_z\rangle_C \tag{158}$$

where $|\uparrow_z\rangle$ and $|\downarrow_z\rangle$ denote the $\sigma_z$-eigenstates. Alice sends one of the particles—the channel particle, denoted by the subscript $C$—to Bob and keeps the ancilla, denoted by $A$. Bob measures either $X, Y$, or $Z$ on the channel particle and returns the channel particle to Alice. Alice then measures an observable $R$ on the pair of particles, where $R$ has the eigenstates (the subscripts $A$ and $C$ are suppressed):

$$|r_1\rangle = \frac{1}{\sqrt{2}}|\uparrow_z\rangle|\uparrow_z\rangle + \frac{1}{2}(|\uparrow_z\rangle|\downarrow_z\rangle e^{i\pi/4} + |\downarrow_z\rangle|\uparrow_z\rangle e^{-i\pi/4}) \tag{159}$$

$$|r_2\rangle = \frac{1}{\sqrt{2}}|\uparrow_z\rangle|\uparrow_z\rangle - \frac{1}{2}(|\uparrow_z\rangle|\downarrow_z\rangle e^{i\pi/4} + |\downarrow_z\rangle|\uparrow_z\rangle e^{-i\pi/4}) \tag{160}$$

$$|r_3\rangle = \frac{1}{\sqrt{2}}|\downarrow_z\rangle|\downarrow_z\rangle + \frac{1}{2}(|\uparrow_z\rangle|\downarrow_z\rangle e^{-i\pi/4} + |\downarrow_z\rangle|\uparrow_z\rangle e^{i\pi/4}) \tag{161}$$

$$|r_4\rangle = \frac{1}{\sqrt{2}}|\downarrow_z\rangle|\downarrow_z\rangle - \frac{1}{2}(|\uparrow_z\rangle|\downarrow_z\rangle e^{-i\pi/4} + |\downarrow_z\rangle|\uparrow_z\rangle e^{i\pi/4}) \tag{162}$$

Note that:

$$|\text{pre}\rangle = \frac{1}{\sqrt{2}}(|\uparrow_z\rangle|\uparrow_z\rangle + |\downarrow_z\rangle|\downarrow_z\rangle \tag{163}$$

$$= \frac{1}{\sqrt{2}}(|\uparrow_x\rangle|\uparrow_x\rangle + |\downarrow_x\rangle|\downarrow_x\rangle \tag{164}$$

$$= \frac{1}{\sqrt{2}}(|\uparrow_y\rangle|\downarrow_y\rangle + |\downarrow_y\rangle|\uparrow_y\rangle \tag{165}$$

$$= \frac{1}{2}(|r_1\rangle + |r_2\rangle + |r_3\rangle + |r_4\rangle) \tag{166}$$

Alice can now assign values to the outcomes of Bob's spin measurements via the ABL-rule, whether Bob measured $X, Y$, or $Z$, based on the post-selections $|r_1\rangle$, $|r_2\rangle$, $|r_3\rangle$, or $|r_4\rangle$, according to Table 1 (where 0 represents the outcome ↑ and 1 represents the outcome ↓) [Vaidman *et al.*, 1987]:

| | $\sigma_x$ | $\sigma_y$ | $\sigma_z$ |
|---|---|---|---|
| $r_1$ | 0 | 0 | 0 |
| $r_2$ | 1 | 1 | 0 |
| $r_3$ | 0 | 1 | 1 |
| $r_4$ | 1 | 0 | 1 |

Table 1: $\sigma_x$, $\sigma_y$, $\sigma_z$ measurement outcomes correlated with eigenvalues of R

This case can be exploited to enable Alice and Bob to share a private random key in the following way: Alice prepares a certain number of copies (depending on the length of the key and the level of privacy desired) of the Bell state $|\text{pre}\rangle$ in Eq. (158). She sends the channel particles to Bob in sequence and keeps the ancillas. Bob measures $X$ or $Z$ randomly on the channel particles and returns the particles, in sequence, to Alice. Alice then measures the observable $R$ on the ancilla and channel pairs and divides the sequence into two subsequences: the subsequence $S_{14}$ for which she obtained the outcomes $r_1$ or $r_4$, and the subsequence $S_{23}$ for which she obtained the outcomes $r_2$ or $r_3$. The sequence of operations can be implemented on a quantum circuit; see [Metzger, 2000].

To check that the channel particles have not been monitored by Eve, Alice now publicly announces (broadcasts) the indices of the subsequence $S_{23}$. As is evident from Table 1, for this subsequence she can make conditional statements of the form: 'For channel particle $i$, if $X$ was measured, the outcome was 1 (0), and if $Z$ was measured, the outcome was 0 (1),' depending on whether the outcome of her $R$-measurement was $r_2$ or $r_3$. She publicly announces these statements as well. If one of these statements, for some index $i$, does not agree with Bob's records, Eve must have monitored the $i$'th channel particle. (Of course, agreement does not entail that the particle was *not* monitored.)

For suppose Eve measures a different spin component observable than Bob on a channel particle and Alice subsequently obtains one of the eigenvalues $r_2$ or $r_3$ when she measures $R$. Bob's measurement outcome, either 0 or 1, will be compatible with just one of these eigenvalues, assuming no intervention by Eve. But after Eve's measurement, both of these eigenvalues will be possible outcomes of Alice's measurement. So Alice's retrodictions of Bob's measurement outcomes for the subsequence $S_{23}$ will not necessarily correspond to Bob's records. In fact, one can show that if Eve measures $X$ or $Z$ randomly on the channel particles, or if she measures a particular one of the observables $X, Y$, or $Z$ on the channel particles (the same observable on each particle), the probability of detection in the subsequence $S_{23}$ is 3/8.

In the subsequence $S_{14}$, the 0 and 1 outcomes of Bob's measurements correspond

to the outcomes $r_1$ and $r_4$ of Alice's $R$-measurements. If, following their public communication about the subsequence $S_{23}$, Alice and Bob agree that there has been no monitoring of the channel particles by Eve, they use the subsequence $S_{14}$ to define a shared raw key.

Note that even a single disagreement between Alice's retrodictions and Bob's records is sufficient to reveal that the channel particles have been monitored by Eve. This differs from the eavesdropping test in the Ekert protocol. Note also that Eve only has access to the channel particles, not the particle pairs. So no strategy is possible in which Eve replaces all the channel particles with her own particles and entangles the original channel particles, treated as a single system, with an ancilla by some unitary transformation, and then delays any measurements until after Alice and Bob have communicated publicly. There is no way that Eve can ensure agreement between Alice and Bob without having access to the particle pairs, or without information about Bob's measurements.

The key distribution protocol as outlined above solves the key distribution problem but not the key storage problem. If Bob actually makes the random choices, measures $X$ or $Z$, and records definite outcomes for the spin measurements before Alice measures $R$, as required by the protocol, Bob's measurement records—stored as classical information—could in principle be copied by Eve without detection. In that case, Eve would know the raw key (which is contained in this information), following the public communication between Alice and Bob to verify the integrity of the quantum communication channel.

To solve the key storage problem, the protocol is modified in the following way: Instead of actually making the random choice for each channel particle, measuring one of the spin observables, and recording the outcome of the measurement, Bob keeps the random choices and the spin measurements 'at the quantum level' until after Alice announces the indices of the subsequence $S_{23}$ of her $R$ measurements. To do this, Bob enlarges the Hilbert space by entangling the quantum state of the channel particle via a unitary transformation with the states of two ancilla particles that he introduces. One particle is associated with a Hilbert space spanned by two eigenstates, $|d_X\rangle$ and $|d_Z\rangle$, of a choice observable or 'quantum die' observable $D$. The other particle is associated with a Hilbert space spanned by two eigenstates, $|p_\uparrow\rangle$ and $|p_\downarrow\rangle$, of a pointer observable $P$. (See §5.2.2 for a discussion of how to implement the unitary transformation on the enlarged Hilbert space.)

On the modified protocol (assuming the ability to store entangled states indefinitely), Alice and Bob share a large number of copies of an entangled 4-particle state. When they wish to establish a random key of a certain length, Alice measures $R$ on an appropriate number of particle pairs in her possession and announces the indices of the subsequence $S_{23}$. Before Alice announces the indices of the subsequence $S_{23}$, neither Alice nor Bob have stored any classical information. So there is nothing for Eve to copy. After Alice announces the indices of the subsequence $S_{23}$, Bob measures the observables $D$ and $P$ on his ancillas with these indices and announces the eigenvalue $|p_\uparrow\rangle$ or $|p_\downarrow\rangle$ as the outcome of his $X$ or $Z$ measurement, depending on the eigenvalue of $D$. If Alice and Bob decide that there has been no eavesdropping by Eve, Bob measures $C$ and $P$ on his ancillas in the subsequence $S_{14}$. It is easy to see that the ABL-rule applies in this case, just as it applies in the case where Bob actually makes the random

choice and actually records definite outcomes of his $X$ or $Z$ measurements before Alice measures $R$. In fact, if the two cases were not equivalent for Alice—if Alice could tell from her $R$-measurements whether Bob had actually made the random choice and actually performed the spin measurements, or had merely implemented these actions 'at the quantum level'—the difference could be exploited to signal superluminally.

## 5.2 Bit Commitment

### 5.2.1 Some History

In a bit commitment protocol, one party, Alice, supplies an encrypted bit to a second party, Bob. The information available in the encrypted bit should be insufficient for Bob to ascertain the value of the bit, but sufficient, together with further information supplied by Alice at a subsequent stage when she is supposed to reveal the value of the bit, for Bob to be convinced that the protocol does not allow Alice to cheat by encrypting the bit in a way that leaves her free to reveal either 0 or 1 at will.

To illustrate the idea, suppose Alice claims the ability to predict the outcomes of elections. To substantiate her claim without revealing valuable information (perhaps to a potential employer, Bob) she suggests the following demonstration: She proposes to record her prediction about whether a certain candidate will win or lose by writing a 0 (for 'lose') or a 1 (for 'win') on a note a month before the election. She will then lock the note in a safe and hand the safe to Bob, but keep the key. After the election, she will announce the bit she chose and prove that she in fact made the commitment at the earlier time by handing Bob the key. Bob can then open the safe and read the note.

Obviously, the security of this procedure depends on the strength of the safe walls or the ingenuity of the locksmith. More generally, Alice can send (encrypted) information to Bob that guarantees the truth of an exclusive classical disjunction (equivalent to her commitment to a 0 or a 1) only if the information is biased towards one of the alternative disjuncts (because a classical exclusive disjunction is true if and only if one of the disjuncts is true and the other false). No principle of classical mechanics precludes Bob from extracting this information, so the security of a classical bit commitment protocol can only be a matter of computational complexity.

The question is whether there exists a quantum analogue of this procedure that is *unconditionally secure*: provably secure as a matter of physical law (according to quantum theory) against cheating by either Alice or Bob. Note that Bob can cheat if he can obtain *some* information about Alice's commitment before she reveals it (which would give him an advantage in repetitions of the protocol with Alice). Alice can cheat if she can delay actually making a commitment until the final stage when she is required to reveal her commitment, or if she can change her commitment at the final stage with a very low probability of detection.

Bennett and Brassard originally proposed a quantum bit commitment protocol in [1984]. The basic idea was to associate the 0 and 1 commitments with two different mixtures represented by the same density operator. As they showed in the same paper, Alice can cheat by adopting an 'EPR attack' or cheating strategy: she prepares entangled pairs of qubits, keeps one of each pair (the ancilla) and sends the second qubit (the channel particle) to Bob. In this way she can fake sending one of two equivalent

mixtures to Bob and reveal either bit at will at the opening stage by effectively steering Bob's particle into the desired mixture by an appropriate measurement. Bob cannot detect this cheating strategy.

In a later paper, Brassard, Crépeau, Josza, and Langlois [1993] proposed a quantum bit commitment protocol that they claimed to be unconditionally secure. The BCJL scheme was first shown to be insecure by Mayers [1996a]. Subsequently, Mayers [1996b, 1997] and Lo and Chau [1997, 1998] independently showed that the insight of Bennett and Brassard in [1984] can be extended to a proof that a generalized version of the EPR cheating strategy can always be applied, if the Hilbert space is enlarged in a suitable way by introducing additional ancilla particles.

The impossibility of unconditionally secure quantum bit commitment came as something of a surprise to the community of quantum cryptologists and has profound consequences. Indeed, it would not be an exaggeration to say that the significance of the quantum bit commitment theorem for our understanding of quantum mechanics is comparable to Bell's theorem [Bell, 1964]. Brassard and Fuchs have speculated ([Brassard, 2000, Fuchs, 1997, Fuchs, 2000, Fuchs and Jacobs, 2002]) that quantum mechanics can be derived from two postulates about quantum information: the possibility of secure key distribution and the impossibility of secure bit commitment. We shall see in §7 what this means for the foundations of quantum mechanics.

Perhaps because of the simplicity of the proof and the universality of the claim, the quantum bit commitment theorem is continually challenged in the literature, on the basis that the proof does not cover all possible procedures that might be exploited to implement quantum bit commitment (see, e.g., Yuen [2005]). There seems to be a general feeling that the theorem is 'too good to be true' and that there must be a loophole.

In fact, there is no loophole. While Kent [1999a, 1999b] has shown how to implement a secure classical bit commitment protocol by exploiting relativistic signalling constraints in a timed sequence of communications between verifiably separated sites for both Alice and Bob, and Hardy and Kent [2004] and Aharonov, Ta-Shma, Vazirani, and Yao [2005] have investigated the security of 'cheat-sensitive' or 'weak' versions of quantum bit commitment, these results are not in conflict with the quantum bit commitment theorem. In a bit commitment protocol as usually understood, there is a time interval of arbitrary length, where no information is exchanged, between the end of the commitment stage of the protocol and the opening or unveiling stage, when Alice reveals the value of the bit. Kent's ingenious scheme effectively involves a third stage between the commitment state and the unveiling stage, in which information is exchanged between Bob's sites and Alice's sites at regular intervals until one of Alice's sites chooses to unveil the originally committed bit. At this moment of unveiling the protocol is not yet complete, because a further sequence of unveilings is required between Alice's sites and corresponding sites of Bob before Bob has all the information required to verify the commitment at a single site. If a bit commitment protocol is understood to require an arbitrary amount of free time between the end of the commitment stage and the opening stage (in which no step is to be executed in the protocol), then the quantum bit commitment theorem covers protocols that exploit special relativistic

signalling constraints.[10]

### 5.2.2 A Key Observation

The crucial insight underlying the proof of the quantum bit commitment theorem is that any step in a quantum bit commitment protocol that requires Alice or Bob to make a definite choice (whether to perform one of a number of alternative measurements, or whether to implement one of a number of alternative unitary transformations) can always be replaced by an EPR cheating strategy in the generalized sense, assuming that Alice and Bob are both equipped with quantum computers. That is, a classical disjunction over definite possibilities—this operation *or* that operation—can always be replaced by a quantum entanglement and a subsequent measurement (perhaps at a more convenient time for the cheater) in which one of the possibilities becomes definite. Essentially, the classical disjunction is replaced by a quantum disjunction. This cheating strategy cannot be detected. Similarly, a measurement can be 'held at the quantum level' without detection: instead of performing the measurement and obtaining a definite outcome as one of a number of possible outcomes, a suitable unitary transformation can be performed on an enlarged Hilbert space, in which the system is entangled with a 'pointer' ancilla in an appropriate way, and the procedure of obtaining a definite outcome can be delayed. The key point is the possibility of keeping the series of transactions between Alice and Bob at the quantum level by enlarging the Hilbert space, until the final exchange of classical information when Alice reveals her commitment.

Any quantum bit commitment scheme will involve a series of transactions between Alice and Bob, where a certain number, $n$, of quantum systems—the 'channel particles'—are passed between them and subjected to various quantum operations (unitary transformations, measurements, etc.), possibly chosen randomly. These operations can always be replaced, without detection, by entangling a channel particle with one or more ancilla particles that function as 'pointer' particles for measurements or 'die' particles for random choices. In effect, this is the (generalized) EPR cheating strategy.

To illustrate: Suppose, at a certain stage of a quantum bit commitment protocol, that Bob is required to make a random choice between measuring one of two observables, $X$ or $Y$, on each channel particle he receives from Alice. For simplicity, assume that $X$ and $Y$ each have two eigenvalues, $x_1$, $x_2$ and $y_1$, $y_2$. After recording the outcome of the measurement, Bob is required to return the channel particle to Alice. When Alice receives the $i$'th channel particle she sends Bob the next channel particle in the sequence. We may suppose that the measurement outcomes that Bob records form part of the information that enables him to confirm Alice's commitment, once she discloses it (together with further information), so he is not required to report his measurement outcomes to Alice until the final stage of the protocol when she reveals her commitment.

Instead of following the protocol, Bob can construct a device that entangles the input state $|\psi\rangle_C$ of a channel particle with the initial states, $|d_0\rangle_B$ and $|p_0\rangle_B$, of two ancilla particles that he introduces, the first of which functions as a 'quantum die' for

[10] I am indebted to Dominic Mayers for clarifying this point.

the random choice and the second as a 'quantum pointer' for the measurement. It is assumed that Bob's ability to construct such a device—in effect, a special purpose quantum computer—is restricted only by the laws of quantum mechanics.

The entanglement is implemented by a unitary transformation in the following way:[11] Define two unitary transformations, $U_X$ and $U_Y$, that implement the $X$ and $Y$ measurements 'at the quantum level' on the tensor product of the Hilbert space of the channel particle, $\mathcal{H}_C$, and the Hilbert space of Bob's pointer ancilla, $\mathcal{H}_{B_P}$:

$$\begin{aligned} |x_1\rangle_C|p_0\rangle_B &\xrightarrow{U_X} |x_1\rangle_C|p_1\rangle_B \\ |x_2\rangle_C|p_0\rangle_B &\xrightarrow{U_X} |x_2\rangle_C|p_2\rangle_B \end{aligned} \tag{167}$$

and

$$\begin{aligned} |y_1\rangle_C|p_0\rangle_B &\xrightarrow{U_Y} |y_1\rangle_C|p_1\rangle_B \\ |y_2\rangle_C||p_0\rangle_B &\xrightarrow{U_Y} |y_2\rangle_C|p_2\rangle_B \end{aligned} \tag{168}$$

so that

$$|\psi\rangle_C|p_0\rangle_B \xrightarrow{U_X} \langle x_1|\psi\rangle|x_1\rangle_C|p_1\rangle_B + \langle x_2|\psi\rangle|x_2\rangle_C|p_2\rangle_B \tag{169}$$

and

$$|\psi\rangle_C|p_0\rangle_B \xrightarrow{U_Y} \langle y_1|\psi\rangle|y_1\rangle_C|p_1\rangle_B + \langle y_2|\psi\rangle|y_2\rangle_C|p_2\rangle_B \tag{170}$$

The random choice is defined similarly by a unitary transformation $V$ on the tensor product of the Hilbert space of Bob's die ancilla, $\mathcal{H}_{B_D}$, and the Hilbert space $\mathcal{H}_C \otimes \mathcal{H}_{B_P}$. Suppose $|d_X\rangle$ and $|d_Y\rangle$ are two orthogonal states in $\mathcal{H}_{B_D}$ and that $|d_0\rangle = \frac{1}{\sqrt{2}}|d_X\rangle + \frac{1}{\sqrt{2}}|d_Y\rangle$. Then (suppressing the obvious subscripts) $V$ is defined by:

$$\begin{aligned} |d_X\rangle \otimes |\psi\rangle|p_0\rangle &\xrightarrow{V} |d_X\rangle \otimes U_X|\psi\rangle|p_0\rangle \\ |d_Y\rangle \otimes |\psi\rangle|p_0\rangle &\xrightarrow{V} |d_Y\rangle \otimes U_Y|\psi\rangle|p_0\rangle \end{aligned} \tag{171}$$

so that

$$\begin{aligned} &|d_0\rangle \otimes |\psi\rangle|p_0\rangle \xrightarrow{V} \\ &\quad \frac{1}{\sqrt{2}}|d_X\rangle \otimes U_X|\psi\rangle|p_0\rangle + \frac{1}{\sqrt{2}}|d_Y\rangle \otimes U_Y|\psi\rangle|p_0\rangle \end{aligned} \tag{172}$$

where the tensor product symbol has been introduced selectively to indicate that $U_x$ and $U_y$ are defined on $\mathcal{H}_C \otimes \mathcal{H}_{B_P}$.

If Bob were to actually choose the observable $X$ or $Y$ randomly, and actually perform the measurement and obtain a particular eigenvalue, Alice's density operator for the channel particle would be:

$$\begin{aligned} &\frac{1}{2}(|\ \langle x_1|\psi\rangle\ |^2\ |x_1\rangle\langle x_1| + |\ \langle x_2|\psi\rangle\ |^2\ |x_2\rangle\langle x_2|) \\ &\quad + \frac{1}{2}(|\ \langle y_1|\psi\rangle\ |^2\ |y_1\rangle\langle y_1| + |\ \langle y_2|\psi\rangle\ |^2\ |y_2\rangle\langle y_2|) \end{aligned} \tag{173}$$

[11] Note that there is no loss of generality in assuming that the channel particle is in a pure state. If the channel particle is entangled with Alice's ancillas, the device implements the entanglement via the transformation $I \otimes \cdots$, where $I$ is the identity operator in the Hilbert space of Alice's ancillas.

assuming that Alice does not know what observable Bob chose to measure, nor what outcome he obtained. But this is precisely the same density operator generated by tracing over Bob's ancilla particles for the state produced in (172). In other words, the density operator for the channel particle is the same for Alice, whether Bob randomly chooses which observable to measure and actually performs the measurement, or whether he implements an EPR cheating strategy with his two ancillas that produces the transition (172) on the enlarged Hilbert space.

If Bob is required to eventually report what measurement he performed and what outcome he obtained, he can at that stage measure the die ancilla for the eigenstate $|d_X\rangle$ or $|d_Y\rangle$, and then measure the pointer ancilla for the eigenstate $|p_1\rangle$ or $|p_2\rangle$. In effect, if we consider the ensemble of possible outcomes for the two measurements, Bob will have converted the 'improper' mixture generated by tracing over his ancillas to a 'proper' mixture. But the difference between a proper and improper mixture is undetectable by Alice since she has no access to Bob's ancillas, and it is only by measuring the composite system consisting of the channel particle together with Bob's ancillas that Alice could ascertain that the channel particle is entangled with the ancillas.

In fact, if it were possible to distinguish between a proper and improper mixture, it would be possible to signal superluminally: Alice could know instantaneously whether or not Bob performed a measurement on his ancillas by monitoring the channel particles in her possession. Note that it makes no difference whether Bob or Alice measures first, since the measurements are of observables in different Hilbert spaces, which therefore commute.

Clearly, a similar argument applies if Bob is required to choose between alternative unitary operations at some stage of a bit commitment protocol. Perhaps less obviously, an EPR cheating strategy is also possible if Bob is required to perform a measurement or choose between alternative operations on channel particle $i + 1$, conditional on the outcome of a prior measurement on channel particle $i$, or conditional on a prior choice of some operation from among a set of alternative operations. Of course, if Bob is in possession of all the channel particles at the same time, he can perform an entanglement with ancillas on the entire sequence, considered as a single composite system. But even if Bob only has access to one channel particle at a time (which he is required to return to Alice after performing a measurement or other operation before she sends him the next channel particle), he can always entangle channel particle $i + 1$ with the ancillas he used to entangle channel particle $i$.

For example, suppose Bob is presented with two channel particles in sequence. He is supposed to decide randomly whether to measure $X$ or $Y$ on the first particle, perform the measurement, and return the particle to Alice. After Alice receives the first particle, she sends Bob the second particle. If Bob measured $X$ on the first particle and obtained the outcome $x_1$, he is supposed to measure $X$ on the second particle; if he obtained the outcome $x_2$, he is supposed to measure $Y$ on the second particle. If he measured $Y$ on the first particle and obtained the outcome $y_1$, he is supposed to apply the unitary transformation $U_1$ to the second particle; if he obtained the outcome $y_2$, he is supposed to apply the unitary transformation $U_2$. After performing the required operation, he is supposed to return the second particle to Alice.

It would seem at first sight that Bob has to actually perform a measurement on the first channel particle and obtain a particular outcome before he can apply the protocol

to the second particle, given that he only has access to one channel particle at a time, so an EPR cheating strategy is excluded. But this is not so. Bob's strategy is the following: He applies the EPR strategy discussed above for two alternative measurements to the first channel particle. For the second channel particle, he applies the following unitary transformation on the tensor product of the Hilbert spaces of his ancillas and the channel particle, where the state of the second channel particle is denoted by $|\phi\rangle$, and the state of the pointer ancilla for the second channel particle is denoted by $|q_0\rangle$ (a second die particle is not required):

$$\begin{aligned} |d_X\rangle|p_1\rangle|\phi\rangle|q_0\rangle &\stackrel{U_C}{\longrightarrow} |d_X\rangle|p_1\rangle \otimes U_X|\phi\rangle|q_0\rangle \\ |d_X\rangle|p_2\rangle|\phi\rangle|q_0\rangle &\stackrel{U_C}{\longrightarrow} |d_X\rangle|p_2\rangle \otimes U_Y|\phi\rangle|q_0\rangle \\ |d_Y\rangle|p_1\rangle|\phi\rangle|q_0\rangle &\stackrel{U_C}{\longrightarrow} |d_Y\rangle|p_1\rangle \otimes U_1|\phi\rangle|q_0\rangle \\ |d_Y\rangle|p_2\rangle|\phi\rangle|q_0\rangle &\stackrel{U_C}{\longrightarrow} |d_Y\rangle|p_2\rangle \otimes U_2|\phi\rangle|q_0\rangle \end{aligned} \qquad (174)$$

### 5.2.3 Proof of the Quantum Bit Commitment Theorem

Since an EPR cheating strategy can always be applied without detection, the proof of the quantum bit commitment theorem assumes that at the end of the commitment stage the composite system consisting of Alice's ancillas, the $n$ channel particles, and Bob's ancillas will be represented by some composite entangled state $|0\rangle$ or $|1\rangle$, depending on Alice's commitment,[12] on a Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$, where $\mathcal{H}_A$ is the Hilbert space of the particles in Alice's possession at that stage (Alice's ancillas and the channel particles retained by Alice, if any), and $\mathcal{H}_B$ is the Hilbert space of the particles in Bob's possession at that stage (Bob's ancillas and the channel particles retained by Bob, if any).

Now, the density operators $W_B(0)$ and $W_B(1)$, characterizing the information available to Bob for the two alternative commitments, are obtained by tracing the states $|0\rangle$ and $|1\rangle$ over $\mathcal{H}_A$. If these density operators are the same, then Bob will be unable to distinguish the 0-state from the 1-state without further information from Alice. In this case, the protocol is said to be 'concealing.' What the proof establishes, by an application of the biorthogonal decomposition theorem, is that if $W_B(0) = W_B(1)$ then there exists a unitary transformation in $\mathcal{H}_A$ that will transform $|0\rangle$ to $|1\rangle$. That is, if the protocol is 'concealing' then it cannot be 'binding' on Alice: she can always follow the protocol (with appropriate substitutions of an EPR strategy) to establish the state $|0\rangle$. At the final stage when she is required to reveal her commitment, she can choose to reveal the alternative commitment, depending on circumstances, by applying a suitable unitary transformation in her own Hilbert space to transform $|0\rangle$ to $|1\rangle$ without Bob being able to detect this move. So either Bob can cheat by obtaining some information about Alice's choice before she reveals her commitment, or Alice can cheat.

The essentials of the proof can be sketched as follows: In the Schmidt decomposi-

[12] More precisely, depending on whether Alice intends to reveal 0 or 1—since we are assuming that Alice will apply an EPR cheating strategy whenever this is relevant.

tion, the states $|0\rangle$ and $|1\rangle$ can be expressed as:

$$\begin{aligned} |0\rangle &= \sum_i \sqrt{p_i}|a_i\rangle|b_i\rangle \\ |1\rangle &= \sum_j \sqrt{p'_j}|a'_j\rangle|b'_j\rangle \end{aligned} \tag{175}$$

where $\{|a_i\rangle\}, \{|a'_j\rangle\}$ are two orthonormal sets of states in $\mathcal{H}_A$, and $\{|b_i\rangle\}, \{|b'_j\rangle\}$ are two orthonormal sets in $\mathcal{H}_B$.

The density operators $W_B(0)$ and $W_B(1)$ are defined by:

$$\begin{aligned} W_B(0) = Tr_A|0\rangle\langle 0| &= \sum_i p_i|b_i\rangle\langle b_i| \\ W_B(1) = Tr_A|1\rangle\langle 1| &= \sum_j p'_j|b'_j\rangle\langle b'_j| \end{aligned} \tag{176}$$

Bob can't cheat if and only if $W_B(0) = W_B(1)$. Now, by the spectral theorem, the decompositions:

$$\begin{aligned} W_B(0) &= \sum_i p_i|b_i\rangle\langle b_i| \\ W_B(1) &= \sum_j p'_j|b'_j\rangle\langle b'_j| \end{aligned}$$

are unique for the nondegenerate case, where the $p_i$ are all distinct and the $p'_j$ are all distinct. The condition $W_B(0) = W_B(1)$ implies that for all $k$:

$$\begin{aligned} p_i &= p'_i \\ |b_i\rangle &= |b'_i\rangle \end{aligned} \tag{177}$$

and so

$$\begin{aligned} |0\rangle &= \sum_i \sqrt{p_i}|a_i\rangle|b_i\rangle \\ |1\rangle &= \sum_i \sqrt{p_i}|a'_i\rangle|b_i\rangle \end{aligned} \tag{178}$$

It follows that there exists a unitary transformation $U \in \mathcal{H}_A$ such that

$$\{|a_k\rangle\} \xrightarrow{U} \{|a'_k\rangle\} \tag{179}$$

and hence

$$|0\rangle \xrightarrow{U} |1\rangle \tag{180}$$

As we shall see in §5.2.4, instead of transforming $|0\rangle$ to $|1\rangle$ by a unitary transformation, Alice could achieve the same effect by preparing the state $|0\rangle$ and measuring in either of two bases, depending on whether she intends to reveal 0 or 1.

The degenerate case can be handled in a similar way. Suppose that $p_1 = p_2 = p'_1 = p'_2 = p$. Then $|b_1\rangle, |b_2\rangle$ and $|b'_1\rangle, |b'_2\rangle$ span the same subspace $\mathcal{H}$ in $\mathcal{H}_B$, and hence (assuming the coefficients are distinct for $k > 2$):

$$\begin{aligned} |0\rangle &= \sqrt{p}(|a_1\rangle|b_1\rangle + |a_2\rangle|b_2\rangle) + \sum_{k>2} \sqrt{p_k}|a_k\rangle|b_k\rangle \\ |1\rangle &= \sqrt{p}(|a'_1\rangle|b'_1\rangle + |a'_2\rangle|b'_2\rangle) + \sum_{k>2} \sqrt{p_k}|a'_k\rangle|b_k\rangle \\ &= \sqrt{p}(|a''_1\rangle|b_1\rangle + |a''_2\rangle|b_2\rangle) + \sum_{k>2} \sqrt{p_k}|a'_k\rangle|b_k\rangle \end{aligned} \quad (181)$$

where $|a''_1\rangle, |a''_2\rangle$ are orthonormal states spanning $\mathcal{H}$. Since $\{|a''_1\rangle, |a''_2\rangle, |a_3\rangle, \ldots\}$ is an orthonormal set in $\mathcal{H}_A$, there exists a unitary transformation in $\mathcal{H}_A$ that transforms $\{|a_k\rangle; k = 1, 2, 3, \ldots\}$ to $\{|a''_1\rangle, |a''_2\rangle, |a'_3\rangle, \ldots\}$, and hence $|0\rangle$ to $|1\rangle$.

The extension of the theorem to the nonideal case, where $W_B(0) \approx W_B(1)$, so that there is a small probability that Bob could distinguish the alternative commitments, shows that Alice has a correspondingly large probability of cheating successfully: there exists a unitary transformation $U$ in $\mathcal{H}_A$ that will transform $W_B(0)$ sufficiently close to $W_B(1)$so that Alice can reveal whichever commitment she chooses, with a corresponding small probability of Bob being able to detect this move.

#### 5.2.4 How the Theorem Works: An Example

The following example by Asher Peres (private communication) is a beautiful illustration of how the theorem works. (My analysis of the example owes much to correspondence with Adrian Kent and Dominic Mayers.)

Suppose Alice is required to send Bob a channel particle $C$ in an equal weight mixture of the qubit states:

$$|c_0\rangle = |0\rangle \quad (182)$$

$$|c_2\rangle = -\frac{1}{2}|0\rangle + \frac{\sqrt{3}}{2}|1\rangle \quad (183)$$

$$|c_4\rangle = -\frac{1}{2}|0\rangle - \frac{\sqrt{3}}{2}|1\rangle \quad (184)$$

if she commits to 0, and an equal weight mixture of the qubit states:

$$|c_1\rangle = |1\rangle \quad (185)$$

$$|c_3\rangle = \frac{\sqrt{3}}{2}|0\rangle - \frac{1}{2}|1\rangle \quad (186)$$

$$|c_5\rangle = -\frac{\sqrt{3}}{2}|0\rangle - \frac{1}{2}|1\rangle \quad (187)$$

if she commits to 1. Note that these two mixtures have the same density operator:

$$\rho_0 = \rho_1 = I/2 \quad (188)$$

Suppose Alice tries to implement an EPR cheating strategy by preparing the entangled state of a system $AC$:

$$|0\rangle = \frac{1}{\sqrt{3}}(|a_0\rangle|c_0\rangle + |a_2\rangle|c_2\rangle + |a_4\rangle|c_4\rangle) \tag{189}$$

where $\{|a_0\rangle, |a_2\rangle, |a_4\rangle\}$ is an orthonormal basis in the 3-dimensional Hilbert space $\mathcal{H}^A$ of a suitable ancilla system $A$. If Alice could transform the state $|0\rangle$ to the state:

$$|1\rangle = \frac{1}{\sqrt{3}}(|a_1\rangle|c_1\rangle + |a_3\rangle|c_3\rangle + |a_5\rangle|c_5\rangle) \tag{190}$$

where $\{|a_1\rangle, |a_3\rangle, |a_5\rangle\}$ is another orthonormal basis in $\mathcal{H}^A$, by a local unitary transformation in $\mathcal{H}^A$, she could delay her commitment to the opening stage. If, at that stage, she decides to commit to 0, she measures the observable with eigenstates $\{|a_0\rangle, |a_2\rangle, |a_4\rangle\}$. If she decides to commit to 1, she performs the local unitary transformation taking the state $|0\rangle$ to the state $|1\rangle$ and measures the observable with eigenstates $\{|a_1\rangle, |a_3\rangle, |a_5\rangle\}$.

Now, $|0\rangle$ can be expressed as:

$$\begin{aligned} |0\rangle &= \frac{1}{\sqrt{3}}\left(|a_0\rangle\frac{|c_3\rangle - |c_5\rangle}{\sqrt{3}} + |a_2\rangle\frac{|c_1\rangle - |c_3\rangle}{\sqrt{3}} + |a_4\rangle\frac{|c_5\rangle - |c_1\rangle}{\sqrt{3}}\right) && (191) \\ &= \frac{1}{\sqrt{3}}\left(\frac{|a_2\rangle - |a_4\rangle}{\sqrt{3}}|c_1\rangle + \frac{|a_0\rangle - |a_2\rangle}{\sqrt{3}}|c_3\rangle + \frac{|a_4\rangle - |a_0\rangle}{\sqrt{3}}|c_5\rangle\right) && (192) \end{aligned}$$

In this representation of $|0\rangle$, the factor states $\frac{|a_2\rangle - |a_4\rangle}{\sqrt{3}}, \frac{|a_0\rangle - |a_2\rangle}{\sqrt{3}}, \frac{|a_4\rangle - |a_0\rangle}{\sqrt{3}}$ in $\mathcal{H}^A$ are not orthogonal—in fact, they are coplanar:

$$|a_0\rangle - |a_2\rangle = -(|a_2\rangle - |a_4\rangle) - (|a_4\rangle - |a_0\rangle) \tag{193}$$

So it seems that there cannot be a suitable unitary transformation that will map $|0\rangle$ to $|1\rangle$ and the EPR strategy is blocked!

Of course, this is not the case. To see that there is such a unitary transformation, note that $|0\rangle$ and $|1\rangle$ can be expressed in the Schmidt decomposition as:

$$\begin{aligned} |0\rangle &= \frac{1}{\sqrt{2}}\left(\frac{2|a_0\rangle - |a_2\rangle - |a_4\rangle}{\sqrt{6}}|c_0\rangle + \frac{|a_2\rangle - |a_4\rangle}{\sqrt{2}}|c_1\rangle\right) && (194) \\ |1\rangle &= \frac{1}{\sqrt{2}}\left(\frac{|a_3\rangle - |a_5\rangle}{\sqrt{2}}|c_0\rangle + \frac{-2|a_1\rangle + |a_3\rangle + |a_5\rangle}{\sqrt{6}}|c_1\rangle\right) && (195) \end{aligned}$$

Clearly, now, there exists a unitary transformation $U$ in $\mathcal{H}^A$ such that:

$$|0\rangle \xrightarrow{U} |1\rangle \tag{196}$$

It follows that:

$$\{|a_0\rangle, |a_2\rangle, |a_4\rangle\} \xrightarrow{U} \{|a'_0\rangle, |a'_2\rangle, |a'_4\rangle\} \tag{197}$$

where $\{|a'_0\rangle, |a'_2\rangle, |a'_4\rangle\}$ is a basis in $\mathcal{H}^A$, and so

$$|1\rangle = \frac{1}{\sqrt{3}}(|a'_0\rangle|c_0\rangle + |a'_2\rangle|c_2\rangle + |a'_4\rangle|c_4\rangle) \quad (198)$$

$$= \frac{1}{\sqrt{3}}(|a_1\rangle|c_1\rangle + |a_3\rangle|c_3\rangle + |a_5\rangle|c_5\rangle) \quad (199)$$

So Alice could implement the EPR cheating strategy by preparing the state $|1\rangle$ and measuring in the basis $\{|a'_0\rangle, |a'_2\rangle, |a'_4\rangle\}$ for the 0-commitment, or in the basis $\{|a_1\rangle, |a_3\rangle, |a_5\rangle\}$ for the 1-commitment. Equivalently, of course, she could prepare the state $|0\rangle$ and measure in two different bases, since the unitary transformation that takes $|1\rangle$ to $|0\rangle$ also takes the basis $\{|a_1\rangle, |a_3\rangle, |a_5\rangle\}$ to the basis $\{|a''_1\rangle, |a''_3\rangle, |a''_5\rangle\}$, and so:

$$|0\rangle = \frac{1}{\sqrt{3}}(|a_0\rangle|c_0\rangle + |a_2\rangle|c_2\rangle + |a_4\rangle|c_4\rangle) \quad (200)$$

$$= \frac{1}{\sqrt{3}}(|a''_1\rangle|c_1\rangle + |a''_3\rangle|c_3\rangle + |a''_5\rangle|c_5\rangle) \quad (201)$$

A calculation shows that:

$$|a''_1\rangle = \frac{1}{3}\left(|a_0\rangle + (1+\sqrt{3})|a_2\rangle + (1-\sqrt{3})|a_4\rangle\right) \quad (202)$$

$$|a''_3\rangle = \frac{1}{3}\left((1+\sqrt{3})|a_0\rangle + (1-\sqrt{3})|a_2\rangle + |a_4\rangle\right) \quad (203)$$

$$|a''_5\rangle = \frac{1}{3}\left(1-\sqrt{3})|a_0\rangle + |a_2\rangle + (1+\sqrt{3})|a_4\rangle\right) \quad (204)$$

In effect, if Alice prepares the entangled state $|0\rangle$ and measures the ancilla $A$ in the $\{|a_0\rangle, |a_2\rangle, |a_4\rangle\}$ basis, she steers the channel particle into a mixture of nonorthogonal states $\{|c_0\rangle, |c_2\rangle, |c_4\rangle\}$. If she measures in the $\{|a''_1\rangle, |a''_3\rangle, |a''_5\rangle\}$ basis, she steers the channel particle into a mixture of nonorthogonal states $\{|c_1\rangle, |c_3\rangle, |c_5\rangle\}$.

It follows that Alice can implement the EPR cheating strategy without performing any unitary transformation—she simply entangles the channel particle with a suitable ancilla particle and performs one of two measurements at the opening stage, depending on her commitment. This shows that the unitary transformation required by the theorem is not in fact required. If a cheating strategy is possible in which Alice, at the opening stage, either makes a measurement on an entangled state for the 0-commitment, or transforms this entangled state to a different state by a local unitary transformation in her Hilbert space and then makes a measurement on the transformed state for the 1-commitment, then an equally good cheating strategy is available in which Alice prepares one entangled state for both commitments, and measures in two alternative bases at the opening stage, depending on her commitment.

### 5.2.5 A Final Worry Laid to Rest

The heart of the mathematical proof is the Schmidt decomposition theorem. But the essential conceptual insight is the possibility of enlarging the Hilbert space and implementing an EPR strategy without detection.

This raises the following question: Suppose Bob cannot cheat because $W_B(0) = W_B(1)$, so by the theorem there exists a unitary transformation $U$ in $\mathcal{H}_A$ that will transform $|0\rangle$ to $|1\rangle$. Could there be a protocol in which Alice also cannot cheat because, although there exists a suitable unitary transformation $U$, she cannot know what unitary transformation to apply? This is indeed the case, but only if $U$ depends on Bob's operations, which are unknown to Alice. But then Bob would have to actually make a definite choice or obtain a definite outcome in a measurement, and he could always avoid doing so without detection by applying an EPR strategy.

This raises a further question: How do we know that following an EPR strategy is never disadvantageous to the cheater? If so, Bob might choose to avoid an EPR strategy in a certain situation because it would be disadvantageous to him. Could there be a bit commitment protocol where the application of an EPR strategy by Bob at a certain stage of the protocol would give Alice the advantage, rather than Bob, while conforming to the protocol would ensure that neither party could cheat? If there were such a protocol, then Bob would, in effect, be forced to conform to the protocol and avoid the EPR strategy, and unconditionally secure bit commitment would be possible.

In fact, the impossibility of such a protocol follows from the theorem (see [Bub, 2001a]. Suppose there were such a protocol. That is, suppose that if Bob applies an EPR strategy then $W_B(0) = W_B(1)$, so by the theorem there exists a unitary transformation $U$ in Alice's Hilbert space that will transform $|0\rangle$ to $|1\rangle$. Alice must know this $U$ because it is uniquely determined by Bob's deviation from the protocol according to an EPR strategy that keeps all disjunctions at the quantum level as linear superpositions. Suppose also that if, instead, Bob is honest and follows the protocol (so that there is a definite choice for every disjunction over possible operations or possible measurement outcomes), then $W_B(0) = W_B(1)$, but the unitary transformation in Alice's Hilbert space that allows her to transform $|0\rangle$ to $|1\rangle$ depends on Bob's choices or measurement outcomes, which are unknown to Alice.

The point to note is that the information available in Alice's Hilbert space must be the same whether Bob follows the protocol and makes determinate choices and obtains determinate measurement outcomes before Alice applies the unitary transformation $U$ that transforms $|0\rangle$ to $|1\rangle$, or whether he deviates from the protocol via an EPR strategy in which he implements corresponding entanglements with his ancillas to keep choices and measurement outcomes at the quantum level before Alice applies the transformation $U$, and only makes these choices and measurement outcomes definite at the final stage of the protocol by measuring his ancillas. There can be no difference for Alice because Bob's measurements on his ancillas and any measurements or operations that Alice might perform take place in different Hilbert spaces, so the operations commute. If Alice's density operator (obtained by tracing over Bob's ancillas), which characterizes the statistics of measurements that Alice can perform in her part of the universe, were different depending on whether or not Bob actually carried out the required measurements, as opposed to keeping the alternatives at the quantum level by implementing corresponding entanglements with ancillas, then it would be possible to use this difference to signal superluminally. Actual measurements by Bob on his ancillas that selected alternatives in the entanglements as determinate would instantaneously alter the information available in Alice's part of the universe.

It follows that in the hypothetical bit commitment protocol we are considering, the

unitary transformation $U$ in Alice's Hilbert space that transforms $|0\rangle$ to $|1\rangle$ must be the same transformation in the honest scenario as in the cheating scenario. But we are assuming that the transformation in the honest scenario is unknown to Alice and depends on Bob's measurement outcomes, while the transformation in the cheating scenario is unique and known to Alice. So there can be no such protocol: the deviation from the protocol by an EPR strategy can never place Bob in a worse position than following the protocol honestly.

The argument can be put formally in terms of the theorem as follows: The cheating scenario produces one of two alternative pure states $|0\rangle_c$ or $|1\rangle_c$ in $\mathcal{H}_A \otimes \mathcal{H}_B$ ('$c$' for 'cheating strategy). Since the reduced density operators in $\mathcal{H}_B$:

$$\begin{aligned} W_B^{(c)}(0) &= Tr_A|0\rangle\langle 0|_c \\ W_B^{(c)}(1) &= Tr_A|1\rangle\langle 1|_c \end{aligned} \tag{205}$$

are required by assumption to be the same:

$$W_B^{(c)}(0) = W_B^{(c)}(1) \tag{206}$$

the states $|0\rangle_c$ and $|1\rangle_c$ can be expressed in biorthogonal decomposition as:

$$\begin{aligned} |0\rangle_c &= \sum_i \sqrt{p_i}|a_i\rangle\langle b_i| \\ |1\rangle_c &= \sum_i \sqrt{p_i}|a_i'\rangle\langle b_i| \end{aligned} \tag{207}$$

where the reduced density operators in $\mathcal{H}_A$:

$$\begin{aligned} W_A^{(c)}(0) = Tr_B|0\rangle\langle 0|_c &= \sum_i p_i|a_i\rangle\langle a_i| \\ W_A^{(c)}(1) = Tr_B|1\rangle\langle 1|_c &= \sum_i p_i|a_i'\rangle\langle a_i'| \end{aligned} \tag{208}$$

are different:

$$W_A^{(c)}(0) \neq W_A^{(c)}(1) \tag{209}$$

It follows that there exists a unitary operator $U_c \in \mathcal{H}_A$ defined by the spectral representations of $W_A^{(c)}(0)$ and $W_A^{(c)}(1)$:

$$\{|a_i\rangle\} \xrightarrow{U_c} \{|a_i'\rangle\} \tag{210}$$

such that:

$$|0\rangle_c \xrightarrow{U_c} |1\rangle_c \tag{211}$$

The honest scenario produces one of two alternative pure states $|0\rangle_h$ and $|1\rangle_h$ in $\mathcal{H}_A \otimes \mathcal{H}_B$ ('$h$' for 'honest'), where the pair $\{|0\rangle_h, |1\rangle_h\}$ depends on Bob's choices and the outcomes of his measurements.

By assumption, as in the cheating scenario, the reduced density operators $W_B^{(h)}(0)$ and $W_B^{(h)}(1)$ in $\mathcal{H}_B$ are the same:

$$W_B^{(h)}(0) = W_B^{(h)}(1) \tag{212}$$

which entails the existence of a unitary operator $U_h \in \mathcal{H}_A$ such that:

$$|0\rangle_h \xrightarrow{U_h} |1\rangle_h \tag{213}$$

where $U_h$ depends on Bob's choices and measurement outcomes.

Now, the difference between the honest scenario and the cheating scenario is undetectable in $\mathcal{H}_A$, which means that the reduced density operators in $\mathcal{H}_A$ are the same in the honest scenario as in the cheating scenario:

$$\begin{aligned} W_A^{(h)}(0) &= W_A^{(c)}(0) \\ W_A^{(h)}(1) &= W_A^{(c)}(1) \end{aligned} \tag{214}$$

Since $U_h$ is defined by the spectral representations of $W_A^{(h)}(0)$ and $W_A^{(h)}(1)$, it follows that $U_h = U_c$. But we are assuming that $U_h$ depends on Bob's choices and measurement outcomes, while $U_c$ is uniquely defined by Bob's EPR strategy, in which there are no determinate choices or measurement outcomes. Conclusion: there can be no bit commitment protocol in which neither Alice nor Bob can cheat if Bob honestly follows the protocol, but Alice can cheat if Bob deviates from the protocol via an EPR strategy. If neither Bob nor Alice can cheat in the honest scenario, then Bob and not Alice must be able to cheat in the cheating scenario.

A similar argument rules out a protocol in which neither party can cheat if Bob is honest (as above), but if Bob follows an EPR strategy, then $W_B(0) \approx W_B(1)$, so Bob has some probability of cheating successfully, but Alice has a greater probability of cheating successfully than Bob. Again, the unitary transformation $U_c$ that would allow Alice to cheat with a certain probability of success if Bob followed an EPR strategy would also have to allow Alice to cheat successfully if Bob were honest. But the supposition is that Alice cannot cheat if Bob is honest, because the unitary transformation $U_h$ in that case depends on Bob's choices and measurement outcomes, which are unknown to Alice. It follows that there can be no such protocol.

So there is no loophole in the proof of the theorem. Unconditionally secure quantum bit commitment (in the sense of the theorem) really is impossible.

# 6 Quantum Computation

## 6.1 The Church-Turing Thesis and Computational Complexity

The classical theory of computation concerns the question of what can be computed, and how efficiently.

Various formal notions of computability by Alonzo Church, Kurt Gödel, and others can all be shown to be equivalent to Alan Turing's notion of computability by a Turing

machine (see, e.g., [Lewis and Papadimitriou, 1981]). A Turing machine is an abstract computational device that can be in one of a finite set of possible states. It has a potentially infinite tape of consecutive cells to store information (0, 1, or blank in each cell) and a movable tape head that reads the information in a cell. Depending on the symbol in a cell and the state of the machine, the tape overwrites the symbol, changes the state, and moves one cell to the right or the left until it finally halts at the completion of the computation. A program for a Turing machine $T$ (e.g., a program that executes a particular algorithm for finding the prime factors of an integer) is a finite string of symbols—which can be expressed as a binary number $b(T)$—indicating, for each state and each symbol, a new state, new symbol, and head displacement. Turing showed that there exists a *universal Turing machine* $U$ that can simulate the program of any Turing machine $T$ with at most a polynomial slow-down, i.e., if we initialize $U$ with $b(T)$ and the input to $T$, then $U$ performs the same computation as $T$, where the number of steps taken by $U$ to simulate each step of $T$ is a polynomial function of $b(T)$. The Church-Turing thesis is the proposal to identify the class of computable functions with the class of functions computable by a universal Turing machine. Equivalently, one could formulate the Church-Turing thesis in terms of decision problems, which have yes-or-no answers (e.g., the problem of determining whether a given number is a prime number).

Intuitively, some computations are harder than others, and some algorithms take more time than others. The computational complexity of an algorithm is measured by the number of steps required by a Turing machine to run through the algorithm. A decision problem is said to be in complexity class **P**, hence *easy* or *tractable* if there exists an algorithm for solving the problem in polynomial time, i.e., in a number of steps that is a polynomial function of the size $n$ of the input (the number of bits required to store the input). A problem is said to be *hard* or *intractable* if there does not exist a polynomial-time algorithm for solving the problem. A problem is in complexity class **EXP** if the most efficient algorithm requires a number of steps that is an exponential function of the size $n$ of the input. The number of steps here refers to the worst-case running time, $\tau$, which is of the order $\mathcal{O}(n^k)$ for a polynomial-time algorithm and of the order $\mathcal{O}(2^n)$ for an exponential-time algorithm.

Note that an exponential-time algorithm could be more efficient than a polynomial-time algorithm for some range of input sizes, so the above terminology should be understood with caution. Consider the following example (taken from [Barenco, 1998, p. 145]): $\tau_P(n) = 10^{-23}n^{1000} + 10^{23}/n \approx \mathcal{O}(n^{1000})$ because, for sufficiently large $n$, the polynomial term dominates (i.e., $\tau_P(n) < cn^{1000}$ for a fixed factor $c$), and $\tau_E(n) = 10^{23}n^{1000} + 10^{-23}2^n \approx \mathcal{O}(2^n)$ because, for sufficiently large $n$, the exponential term dominates (i.e., $\tau_E(n) < c2^n$ for a fixed factor $c$). But for small enough values of $n$, $\tau_E(n) < \tau_P(n)$.

A Turing machine as defined above is a deterministic machine. A *nondeterministic* or *probabilistic* Turing machine makes a random choice between multiple transitions (to a new symbol, new state, and head displacement) for each symbol and each state. For each sequence of choices, the sequence of transitions corresponds to a sequence of steps executed by a deterministic Turing machine. If any of these machines halts, the computation is regarded as completed. Evidently, a nondeterministic Turing machine cannot compute a function that is not computable by a deterministic Turing machine,

but it is believed (but not proved) that certain problems can be solved more efficiently by nondeterministic Turing machines than by any deterministic Turing machine. The complexity class **NP** is the class of problems that can be solved in polynomial time by a nondeterministic Turing machine. This is equivalent to the class of problems for which proposed solutions can be verified in polynomial time by a deterministic Turing machine. For example, it is believed (but not proved) that the problem of factoring an integer into its prime factors is a 'hard' problem: there is no known polynomial-time algorithm for this problem. However, the problem of checking whether a candidate factor of an integer is indeed a factor can be solved in polynomial time, so factorizability is an **NP** problem.

Clearly **P** $\subseteq$ **NP**, but it is an open problem in complexity theory whether **P** = **NP**. An **NP** problem is said to be **NP**-complete if every **NP** problem can be shown to have a solution with a number of steps that is a polynomial function of the number of steps required to solve the **NP**-complete problem. So if an **NP**-complete problem can be solved in polynomial time, then all **NP** problems can be solved in polynomial time, and **P** = **NP**. The problem of determining whether a Boolean function $f\{0,1\}^n \rightarrow \{0,1\}$ is satisfiable (i.e., whether there is a set of input values for which the function takes the value 1, or equivalently whether there is an assignment of truth values to the atomic sentences of a compound sentence of Boolean logic under which the compound sentence comes out true) is an **NP**-complete problem. Factorizability is an **NP** problem that is not known to be **NP**-complete.

Since a Turing machine can simulate any classical computing device with at most a polynomial slow-down, the complexity classes are the same for any model of computation. For example, a circuit computer computes the value of a function by transforming data stored in an input register, representing the input to the function, via Boolean circuits constructed of elementary Boolean gates connected by wires, to data in an output register representing the value of the function computed. The elementary Boolean gates are 1-bit gates (such as the NOT gate, which transforms 0 to 1, and conversely) and 2-bit gates (such as the AND gate, which takes two input bits to 1 if and only if they are both 1, otherwise to 0), and it can be shown that a combination of such gates forms a 'universal set' that suffices for any transformation of $n$ bits. In fact, it turns out that one of the sixteen possible 2-bit Boolean gates, the NAND gate (or NOT AND gate), which takes two input bits to 0 of and only if they are both 1, forms a universal set by itself.

In a circuit model of a quantum computer, the registers store qubits, which are then manipulated by elementary unitary gates. It can be shown (see [Nielsen and Chuang, 2000, p. 188]) that a set of single-qubit and two-qubit unitary gates—the CNOT gate, the Hadamard gate, the phase gate, and the $\pi/8$ gate—forms a universal set, in the sense that any unitary transformation of $n$ qubits can be approximated to arbitrary accuracy by a quantum circuit consisting of these gates connected in some combination. The CNOT gate ('C' for 'controlled) has two input qubits, a 'control' qubit and a 'target' qubit. The gate functions so as to flip the target qubit if and only if the control qubit is $|1\rangle$. The remaining three gates are single-qubit gates. The Hadamard gate transforms $|0\rangle$ to $(|0\rangle + |1\rangle)/\sqrt{2}$ and $|1\rangle$ to $(|0\rangle - |1\rangle)/\sqrt{2}$ and is sometimes referred to as the 'square root of NOT' gate because two successive applications transforms $|0\rangle$ to $|1\rangle$, and conversely. The phase gate leaves $|0\rangle$ unchanged and trans-

forms $|1\rangle$ to $i|1\rangle$. The $\pi/8$ gate leaves $|0\rangle$ unchanged and transforms $|1\rangle$ to $e^{i\pi/4}|1\rangle$. (See [Nielsen and Chuang, 2000, p. 174] for a discussion and why the $\pi/8$ gate is so named.)

There are other models of quantum computation. In the 'cluster state' or 'one-way' quantum computer of Raussendorf and Briegel [2001b, 2001a], a fixed multi-qubit entangled state (called a 'cluster state'), independent of the computation, is prepared. Then a sequence of single-qubit measurements is performed on this state, where the choice of what observables to measure depends on the outcomes of the previous measurements. No unitary transformations are involved. Remarkably, it can be shown that any quantum circuit of unitary gates and measurements can be simulated by a cluster state computer with similar resources of qubits and time (see [Jozsa, 2005, Nielsen, 2003, Nielsen, 2005]).

The interesting question is whether a quantum computer can perform computational tasks that are not possible for a Turing machine, or perform such tasks more efficiently than any Turing machine. Since a Turing machine is defined by its program, and a program can be specified by a finite set of symbols, there are only countably many Turing machines. There are uncountably many functions on the natural numbers, so there are uncountably many uncomputable functions, i.e., functions that are not computable by any Turing machine. A quantum computer cannot compute a function that is not Turing-computable, because a Turing machine can simulate (albeit inefficiently, with an exponential slow-down [Feynman, 1982]) the dynamical evolution of any system, classical or quantum, with arbitrary accuracy. But there are computational tasks that a quantum computer can perform by exploiting entanglement that are impossible for any Turing machine. Recall the discussion of Bell's counterargument to the EPR argument in §3.1.1: a quantum computer, but no classical computer, can perform the task of rapidly producing pairs of values (0 or 1) for pairs of input angles at different locations, with correlations that violate Bell's inequality, where the response time is less than the time taken by light to travel between the locations.

The current interest in quantum computers concerns the question of whether a quantum computer can compute certain Turing-computable functions more *efficiently* than any Turing machine. In the following section, I discuss quantum algorithms that achieve an exponential speed-up over any classical algorithm, or an exponential speed-up over any *known* classical algorithm. The most spectacular of these is Shor's factorization algorithm, and a related algorithm for solving the discrete log problem.[13]

The factorization algorithm has an important practical application to cryptography. Public-key distribution protocols such as RSA [Rivest *et al.*, 1978] (widely used in commercial transactions over the internet, transactions between banks and financial institutions, etc.) rely on factoring being a 'hard' problem. (Preskill [2005] notes that currently the 65-digit factors of a 130-digit integer can be found in about a month using a network of hundreds of work stations implementing the best known classical factoring algorithm (the 'number sieve algorithm'). He estimates that factoring a 400-digit integer would take about $10^{10}$ years, which is the age of the universe.) To see the idea behind the RSA protocol, suppose Alice wishes to send a secret message to Bob. Bob's

[13]The discrete log of $x$ with respect to a given prime integer $p$ and an integer $q$ coprime to $p$ is the integer $r$ such that $q^r = x \bmod p$. See [Nielsen and Chuang, 2000, p. 238] for a discussion.

public key consists of two large integers, $s$ and $c$. Alice encrypts the message $m$ (in the form of a binary number) as $e = m^s \bmod c$ and sends the encrypted message to Bob. Bob decrypts the message as $e^t \bmod c$ where $t$ is an integer known only to Bob. The integer $t$ for which $m = e^t \bmod c$ can easily be determined from $s$ and the factors of $c$, but since $c = pq$ is the product of two large prime numbers known only to Bob, an eavesdropper, Eve, can read the message only if she can factor $c$ into its prime factors. The cleverness of the scheme resides in the fact that no secret key needs to be distributed between Alice and Bob: Bob's key $\{s, c\}$ is public and allows anyone to send encrypted messages to Bob. If a quantum computer could be constructed that implemented Shor's algorithm, key distribution protocols that rely on the difficulty of factoring very large numbers would be insecure.

## 6.2 Quantum Algorithms

In the following three sections, I look at the information-processing involved in Deutsch's XOR algorithm [1985], Simon's period-finding algorithm [1994, 1997], and Shor's factorization algorithm [1994, 1997] in terms of the difference between the Boolean logic underlying a classical computation and the non-Boolean logic represented by the projective geometry of Hilbert space, in which the subspace structure of Hilbert space replaces the set-theoretic structure of classical logic. The three algorithms all turn out to involve a similar geometric formulation.

Basically, all three algorithms involve the determination of a global property of a function, i.e., a disjunctive property. The disjunction is represented as a subspace in an appropriate Hilbert space, and alternative possible disjunctions turn out to be represented as orthogonal subspaces, except for intersections or overlaps. The true disjunction is determined as the subspace containing the state vector via a measurement. The algorithm generally has to be run several times because the state might be found in the overlap region. The essential feature of these quantum computations is that the true disjunction is distinguished from alternative disjunctions without determining the truth values of the disjuncts. In a classical computation, distinguishing the true disjunction would be impossible without the prior determination of the truth values of the disjuncts. More generally, a quantum computer computes a global property of a function without computing information that is redundant quantum mechanically, but essential for a classical computation of the global property.

There are other quantum algorithms besides these three, e.g., Grover's sorting algorithm [1997] which achieves a quadratic speed-up over any classical algorithm. For a discussion, see [Nielsen and Chuang, 2000], [Jozsa, 1999].

### 6.2.1 Deutsch's XOR Algorithm and the Deutsch-Jozsa Algorithm

Let $B = \{0, 1\}$ be a Boolean algebra (or the additive group of integers mod 2). In Deutsch's XOR problem [1985], we are given a 'black box' or oracle that computes a function $f : B \rightarrow B$ and we are required to determine whether the function is 'constant' (takes the same value for both inputs) or ''balanced' (takes a different value for each input). Classically, the only way to do this would be to consult the oracle twice, for the input values 0 and 1, and compare the outputs.

In a quantum computation of the Boolean function, a unitary transformation $U_f$ : $|x\rangle|y\rangle \rightarrow |x\rangle|y \oplus f(x)\rangle$ corresponding to the 'black box' correlates input values with corresponding output values.[14] The computation proceeds as follows: The input and output registers are 1-qubit registers initialized to the state $|0\rangle|0\rangle$ in a standard basis. A Hadamard transformation is applied to the input register, yielding a linear superposition of states corresponding to the two possible input values 0 and 1, and the transformation $U_f$ is then applied to both registers, yielding the transitions:

$$|0\rangle|0\rangle \stackrel{H}{\rightarrow} \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)|0\rangle \qquad (215)$$

$$\stackrel{U_f}{\rightarrow} \frac{1}{\sqrt{2}}(|0\rangle|f(0)\rangle + |1\rangle|f(1)\rangle) \qquad (216)$$

If the function is constant, the final composite state of both registers is one of the two orthogonal states:

$$|c_1\rangle = \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|0\rangle) \qquad (217)$$

$$|c_2\rangle = \frac{1}{\sqrt{2}}(|0\rangle|1\rangle + |1\rangle|1\rangle) \qquad (218)$$

If the function is balanced, the final composite state is one of the two orthogonal states:

$$|b_1\rangle = \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|1\rangle) \qquad (219)$$

$$|b_2\rangle = \frac{1}{\sqrt{2}}(|0\rangle|1\rangle + |1\rangle|0\rangle) \qquad (220)$$

The states $|c_1\rangle, |c_2\rangle$ and $|b_1\rangle, |b_2\rangle$ span two planes $P_c, P_b$ in $\mathcal{H}^2 \otimes \mathcal{H}^2$, represented by the projection operators:

$$P_c = P_{|c_1\rangle} + P_{|c_2\rangle} \qquad (221)$$

$$P_b = P_{|b_1\rangle} + P_{|b_2\rangle} \qquad (222)$$

These planes are orthogonal, except for an intersection, so their projection operators commute. The intersection is the line (ray) spanned by the vector[15]:

$$\frac{1}{2}(|00\rangle + |01\rangle + |10\rangle + |11\rangle) = \frac{1}{\sqrt{2}}(|c_1\rangle + |c_2\rangle) = \frac{1}{\sqrt{2}}(|b_1\rangle + |b_2\rangle) \qquad (223)$$

In the 'prime' basis spanned by the states $|0'\rangle = H|0\rangle, |1'\rangle = H|1\rangle$ the intersection is the state $|0'\rangle|0'\rangle$, the 'constant' plane is spanned by $|0'\rangle|0'\rangle, |0'\rangle|1'\rangle$, and the

[14] Note that two quantum registers are required to compute functions that are not 1-1 by a unitary transformation. Different input values, $x$ and $y$, to a function $f$ are represented by orthogonal states $|x\rangle, |y\rangle$. So if $f(x) = f(y)$ for some $x \neq y$, the transformation $W_f : |x\rangle \rightarrow |f(x)\rangle$ could not be unitary, because the orthogonal states $|x\rangle, |y\rangle$ would have to be mapped onto the same state by $W_f$. The ability of unitary transformations, which are reversible, to compute irreversible functions is achieved by keeping a record of the input for each output value of the function.

[15] Here $|00\rangle = |0\rangle|0\rangle$, etc.

'balanced' plane is spanned by $|0'\rangle|0'\rangle, |1'\rangle|1'\rangle$. Note that:

$$|0'\rangle|1'\rangle = \frac{1}{\sqrt{2}}(|c_1\rangle - |c_2\rangle) \tag{224}$$

$$|1'\rangle|1'\rangle = \frac{1}{\sqrt{2}}(|b_1\rangle - |b_2\rangle) \tag{225}$$

In the usual formulation of the algorithm, to decide whether the function $f$ is constant or balanced we measure the output register in the prime basis. If the outcome is $0'$ (which is obtained with probability 1/2, whether the state ends up in the constant plane or the balanced plane), the computation is inconclusive, yielding no information about the function $f$. If the outcome is $1'$, then we measure the input register. If the outcome of the measurement on the input register is $0'$, the function is constant; if it is $1'$, the function is balanced.

Alternatively—and this will be relevant for the comparison with Simon's algorithm and Shor's algorithm—we could measure the observable with eigenstates $|0'0'\rangle$, $|0'1'\rangle$, $|1'0'\rangle$, $|1'1'\rangle$. The final state is in the 3-dimensional subspace orthogonal to the vector $|1'0'\rangle$, either in the constant plane or the balanced plane. If the state is in the constant plane, we will either obtain the outcome $0'0'$ with probability 1/2 (since the final state is at an angle $\pi/4$ to $|0'0'\rangle$), in which case the computation is inconclusive, or the outcome $0'1'$ with probability 1/2. If the state is in the balanced plane, we will again obtain the outcome $0'0'$ with probability 1/2, in which case the computation is inconclusive, or the outcome $1'1'$ with probability 1/2. So in either case, with probability 1/2, we can distinguish in one run of the algorithm between the two quantum disjunctions 'constant' and 'balanced' represented by the planes:

$$P_c = P_{|0'0'\rangle} \vee P_{|0'1'\rangle} \tag{226}$$

$$P_b = P_{|0'0'\rangle} \vee P_{|1'1'\rangle} \tag{227}$$

without finding out the truth values of the disjuncts in the computation (i.e., whether in the 'constant' case the function maps 0 to 0 and 1 to 0 or whether the function maps 0 to 1 and 1 to 1, and similarly in the 'balanced' case). Note that we could also apply a Hadamard transformation to the final states of both registers and measure in the computational basis, since $|0'0'\rangle \xrightarrow{H} |00\rangle$, etc.

Deutsch's XOR algorithm was the first quantum algorithm with a demonstrated speed-up over any classical algorithm performing the same computational task. However, the algorithm has an even probability of failing, so the improvement in efficiency over a classical computation is only achieved if the algorithm succeeds, and even then is rather modest: one run of the quantum algorithm versus two runs of a classical algorithm. The following variation of Deutsch's algorithm avoids this feature [Cleve *et al.*, 1998].

We begin by initializing the two registers to $|0\rangle$ and $|1\rangle$, respectively (instead of to $|0\rangle$ and $|0\rangle$) and apply a Hadamard transformation to both registers, which yields the transition:

$$|0\rangle|1\rangle \xrightarrow{H} \frac{|0\rangle + |1\rangle}{\sqrt{2}}\frac{|0\rangle - |1\rangle}{\sqrt{2}} \tag{228}$$

Since

$$U_f|x\rangle|y\rangle = |x\rangle|y \oplus f(x)\rangle \tag{229}$$

it follows that

$$U_f|x\rangle\frac{|0\rangle - |1\rangle}{\sqrt{2}} = \begin{cases} |x\rangle\frac{|0\rangle-|1\rangle}{\sqrt{2}} \text{ if } f(x) = 0 \\ -|x\rangle\frac{|0\rangle-|1\rangle}{\sqrt{2}} \text{ if } f(x) = 1 \end{cases} \tag{230}$$

which can be expressed as

$$U_f|x\rangle\frac{|0\rangle - |1\rangle}{\sqrt{2}} = (-1)^{f(x)}|x\rangle\frac{|0\rangle - |1\rangle}{\sqrt{2}} \tag{231}$$

Notice that the value of the function now appears as a phase of the final state of the input register, a feature referred to as 'phase kickback.' For the input state $1/\sqrt{2}(|0\rangle + |1\rangle)$, we have:

$$U_f\frac{|0\rangle + |1\rangle}{\sqrt{2}}\frac{|0\rangle - |1\rangle}{\sqrt{2}} = \frac{(-1)^{f(0)}|0\rangle + (-1)^{f(1)}|1\rangle}{\sqrt{2}}\frac{|0\rangle - |1\rangle}{\sqrt{2}} \tag{232}$$

which can be expressed as:

$$U_f\frac{|0\rangle + |1\rangle}{\sqrt{2}}\frac{|0\rangle - |1\rangle}{\sqrt{2}} = \begin{cases} \pm\frac{|0\rangle+|1\rangle}{\sqrt{2}}\frac{|0\rangle-|1\rangle}{\sqrt{2}} = \pm|0'\rangle|1'\rangle \text{ if } f(0) = f(1) \\ \pm\frac{|0\rangle-|1\rangle}{\sqrt{2}}\frac{|0\rangle-|1\rangle}{\sqrt{2}} = \pm|1'\rangle|1'\rangle \text{ if } f(0) \neq f(1) \end{cases} \tag{233}$$

Instead of the final state of the two registers ending up as one of two orthogonal states in the constant plane, or as one of two orthogonal states in the balanced plane, the final state now ends up as $\pm|0'1'\rangle$ in the constant plane, or as $\pm|1'1'\rangle$ in the balanced plane, and these states can be distinguished because they are orthogonal. So we can decide with certainty whether the function is constant or balanced after only one run of the algorithm. In fact, we can distinguish these two possibilities by simply measuring the input register in the prime basis. Note that if we perform a final Hadamard transformation on the input register (which takes $|0'\rangle$ to $|0\rangle$ and $|1'\rangle$ to $|1\rangle$), we can distinguish the two possibilities by measuring the input register in the computational basis. Note also that the state of the output register is unchanged: at the end of the process it is in the state $|1'\rangle = H|1\rangle$ (as in (228)) and is not measured.

Deutsch's XOR problem can be generalized to the problem ('Deutsch's problem') of determining whether a Boolean function $f : B^n \to B$ is constant or whether it is balanced, where it is promised that the function is either constant or balanced. 'Balanced' here means that the function takes the values 0 and 1 an equal number of times, i.e., $2^{n-1}$ times each. The Deutsch-Jozsa algorithm [1992] decides whether $f$ is constant or balanced in one run.

We begin by setting the input $n$-qubit register to the state $|0\rangle$ (an abbreviation for the state $|0 \cdots 0\rangle = |0\rangle \cdots |0\rangle$) and the output 1-qubit register to the state $|1\rangle$. We apply an $n$-fold Hadamard transformation to the input register and a Hadamard transformation to the output register, followed by the unitary transformation $U_f$ to both registers, and finally an $n$-fold Hadamard transformation to the input register.

First note that

$$H|x\rangle = \frac{1}{\sqrt{2}} \sum_{y \in \{0,1\}} (-1)^{xy} |y\rangle \tag{234}$$

so

$$H^{\otimes n}|x_1, \ldots, x_n\rangle = \frac{1}{\sqrt{2}} \sum_{y_1,\cdots,y_n} (-1)^{x_1 y_1 + \cdots + x_n y_n} |y_1, \ldots, y_n\rangle \tag{235}$$

This can be expressed as:

$$H^{\otimes n}|x\rangle = \frac{1}{\sqrt{2}} \sum_{y \in \{0,1\}} (-1)^{x \cdot y} |y\rangle \tag{236}$$

where $x \cdot y$ is the bitwise inner product of $x$ and $y$, mod 2.

The unitary transformations (Hadamard transformation, $U_f$) yield:

$$|0\rangle^{\otimes n}|1\rangle \quad \xrightarrow{H} \quad \sum_{x \in \{0,1\}^n} \frac{|x\rangle}{\sqrt{2^n}} \frac{|0\rangle - |1\rangle}{\sqrt{2}} \tag{237}$$

$$\xrightarrow{U_f} \quad \sum_x \frac{(-1)^{f(x)}}{\sqrt{2^n}} |x\rangle \frac{|0\rangle - |1\rangle}{\sqrt{2}} \tag{238}$$

$$\xrightarrow{H} \quad \sum_y \sum_x \frac{(-1)^{x \cdot y + f(x)}}{\sqrt{2^n}} |y\rangle \frac{|0\rangle - |1\rangle}{\sqrt{2}} \tag{239}$$

Now consider the state of the input register:

$$\sum_y \sum_x \frac{(-1)^{x \cdot y + f(x)}}{\sqrt{2^n}} |y\rangle = \sum_x \frac{(-1)^{f(x)}}{\sqrt{2^n}} |0 \ldots 0\rangle + \ldots \tag{240}$$

Note that the coefficient (amplitude) of the state $|0 \ldots 0\rangle$ in the linear superposition (240) is $\sum_x \frac{(-1)^{f(x)}}{\sqrt{2^n}}$. If $f$ is constant, this coefficient is $\pm 1$, so the coefficients of the other terms must all be 0. If $f$ is balanced, $f(x) = 0$ for half the values of $x$ and $f(x) = 1$ for the other half, so the positive and negative contributions to the coefficient of $|0 \ldots 0\rangle$ cancel to 0. In other words, if $f$ is constant, the state of the input register is $\pm|0 \ldots 0\rangle$; if $f$ is balanced, the state is in the orthogonal subspace.

This is the usual way of describing how the algorithm works, which rather obscures the geometric picture. Consider, for simplicity, the case $n = 2$. After the transformation $U_f$, but before the final Hadamard transformation, the state of the input register is:

$$\pm \frac{1}{2}(|00\rangle + |01\rangle + |10\rangle + |11\rangle) \tag{241}$$

if the function is constant, or a state of the form:

$$\frac{1}{2}(\pm|00\rangle \pm |01\rangle \pm |10\rangle \pm |11\rangle) \tag{242}$$

if the function is balanced, where two of the coefficients are $+1$ and two of the coefficients are $-1$. Evidently, there are six such balanced states, and they are all orthogonal

to the constant state. So the six balanced states lie in a 3-dimensional subspace orthogonal to the constant state and can therefore be distinguished from the constant state. The final Hadamard transformation transforms the constant state:

$$\pm\frac{1}{2}(|00\rangle + |01\rangle + |10\rangle + |11\rangle) \xrightarrow{H} \pm|00\rangle \tag{243}$$

and the six balanced states to states in the 3-dimensional subspace orthogonal to $|00\rangle$. So to decide whether the function is constant or balanced we need only measure the input register and check whether it is in the state $|00\rangle$.

### 6.2.2 Simon's Algorithm

The problem here is to find the period $r$ of a periodic function $f : B^n \to B^n$, i.e., a Boolean function for which

$$f(x_i) = f(x_j) \text{ if and only if } x_j = x_i \oplus r, \text{ for all } x_i, x_j \in B^n \tag{244}$$

Note that since $x \oplus r \oplus r = x$, the function is 2-to-1.

Simon's algorithm solves the problem efficiently, with an exponential speed-up over any classical algorithm (see [Simon, 1994, Simon, 1997]). The algorithm proceeds as in the Deutsch-Jozsa algorithm, starting with the input and output registers in the state $|0\dots0\rangle|0\rangle$ in the computational basis:

$$|0\dots0\rangle|0\rangle \xrightarrow{H} \frac{1}{\sqrt{2^n}}\sum_{x=0}^{2^n-1}|x\rangle|0\rangle \tag{245}$$

$$\xrightarrow{U_f} \frac{1}{\sqrt{2^n}}\sum_x |x\rangle|f(x)\rangle \tag{246}$$

$$= \frac{1}{\sqrt{2^{n-1}}}\sum_{x_i}\frac{|x_i\rangle + |x_i \oplus r\rangle}{\sqrt{2}}|f(x_i)\rangle \tag{247}$$

where $U_f$ is the unitary transformation implementing the Boolean function as:

$$U_f : |x\rangle|y\rangle \to |x\rangle|y \oplus f(x)\rangle \tag{248}$$

The usual way to see how the algorithm works is to consider what happens if we measure the output register and keep the state of the input register,[16] which will have the form:

$$\frac{|x_i\rangle + |x_i \oplus r\rangle}{\sqrt{2}} \tag{249}$$

This state contains the information $r$, but summed with an unwanted randomly chosen offset $x_i$ that depends on the measurement outcome. A direct measurement of the state label would yield any $x \in B^n$ equiprobably, providing no information about $r$.

[16] The measurement of the output register here is a pedagogical device for ease of conceptualization. Only the input register is actually measured. The input register is in a mixture of states, which we can think of as the mixture associated with the distribution of outcomes obtained by measuring the output register.

We now apply a Hadamard transform:

$$\frac{|x_i\rangle + |x_i \oplus r\rangle}{\sqrt{2}} \quad \xrightarrow{H} \quad \sum_{y \in B^n} \frac{(-1)^{x_i \cdot y} + (-1)^{(x_i \oplus r)\cdot y}}{\sqrt{2}} |y\rangle \tag{250}$$

$$= \quad \sum_{y : r \cdot y = 0} \frac{(-1)^{x_i \cdot y}}{\sqrt{2}} |y\rangle \tag{251}$$

where the last equality follows because terms interfere destructively if $r \cdot y = 1$. Finally, we measure the input register in the computational basis and obtain a value $y$ (equiprobably) such that $r \cdot y = 0$. Then we repeat the algorithm sufficiently many times to find enough values $y_i$ so that $r$ can be determined by solving the linear equations $r \cdot y_1 = 0, \ldots, r \cdot y_k = 0$.

To see what is going on geometrically, consider the case $n = 2$. The possible values of the period $r$ are: 01, 10, 11, and the corresponding states of the input and output registers after the unitary transformation $U_f$ are:

$$r = 01 : \; (|00\rangle + |01\rangle)|f(00)\rangle + (|10\rangle + |11\rangle)|f(10)\rangle$$

$$r = 10 : \; (|00\rangle + |10\rangle)|f(00)\rangle + (|01\rangle + |11\rangle)|f(01)\rangle$$

$$r = 11 : \; (|00\rangle + |11\rangle)|f(00)\rangle + (|01\rangle + |10\rangle)|f(01)\rangle$$

Notice that this case reduces to the same geometric construction as in Deutsch's XOR algorithm. For $r = 10$ the input register states are $|c_1\rangle = |00\rangle + |10\rangle$ or $|c_2\rangle = |01\rangle + |11\rangle$, and for $r = 11$ the input register states are $|b_1\rangle = |00\rangle + |11\rangle$ or $|b_2\rangle = |01\rangle + |10\rangle$, depending on the outcome of the measurement of the output register. So the three possible periods are associated with three planes in $\mathcal{H}^2 \otimes \mathcal{H}^2$, which correspond to the constant and balanced planes in Deutsch's XOR algorithm, and a third plane, all three planes intersecting in the line spanned by the vector $|00\rangle$. In the prime basis obtained by applying the Hadamard transformation, the planes are as follows:

$r = 01$ : plane spanned by $|0'0'\rangle, |1'0'\rangle$

$r = 10$ : plane spanned by $|0'0'\rangle, |0'1'\rangle$ (corresponds to 'constant' plane)

$r = 11$ : plane spanned by $|0'0'\rangle, |1'1'\rangle$ (corresponds to 'balanced' plane)

We could simply measure the input register in the prime basis to find the period. Alternatively, we could apply a Hadamard transformation (which amounts to dropping the primes in the above representation of the $r$-planes) and measure in the computational basis. The three planes are orthogonal, except for their intersection in the line spanned by the vector $|00\rangle$. The three possible periods can therefore be distinguished by measuring the observable with eigenstates $|00\rangle, |01\rangle, |10\rangle, |11\rangle$, except when the state of the register is projected by the measurement ('collapses') onto the intersection state $|00\rangle$ (which occurs with probability 1/2). So the algorithm will generally have to be repeated until we find an outcome that is not 00.

The $n = 2$ case of Simon's algorithm reduces to Deutsch's XOR algorithm. What about other cases? We can see what happens in the general case if we consider the case

$n = 3$. There are now seven possible periods: 001, 010, 011, 100, 101, 110, 111. For the period $r = 001$, the state of the two registers after the unitary transformation $U_f$ is:

$$(|000\rangle + |001\rangle)|f(000)\rangle + (|010\rangle + |011\rangle)|f(010)\rangle \\ +(|100\rangle + |101\rangle)|f(100)\rangle + (|110\rangle + |111\rangle)|f(110)\rangle \quad (252)$$

If we measure the output register, the input register is left in one of four states, depending on the outcome of the measurement:

$$\begin{aligned} |000\rangle + |001\rangle &= |0'0'0'\rangle + |0'1'0'\rangle + |1'0'0'\rangle + |1'1'0'\rangle \\ |010\rangle + |011\rangle &= |0'0'0'\rangle - |0'1'0'\rangle + |1'0'0'\rangle - |1'1'0'\rangle \\ |100\rangle + |101\rangle &= |0'0'0'\rangle + |0'1'0'\rangle - |1'0'0'\rangle - |1'1'0'\rangle \\ |110\rangle + |111\rangle &= |0'0'0'\rangle - |0'1'0'\rangle - |1'0'0'\rangle + |1'1'0'\rangle \end{aligned}$$

Applying a Hadamard transformation amounts to dropping the primes. So if the period is $r = 001$, the state of the input register ends up in the 4-dimensional subspace of $\mathcal{H}^2 \otimes \mathcal{H}^2 \otimes \mathcal{H}^2$ spanned by the vectors: $|000\rangle, |010\rangle, |100\rangle, |110\rangle$.

A similar analysis applies to the other six possible periods. The corresponding subspaces are spanned by the following vectors:

r = 001: $|000\rangle, |010\rangle, |100\rangle, |110\rangle$

r = 010: $|000\rangle, |001\rangle, |100\rangle, |101\rangle$

r = 011: $|000\rangle, |011\rangle, |100\rangle, |111\rangle$

r = 100: $|000\rangle, |001\rangle, |010\rangle, |011\rangle$

r = 101: $|000\rangle, |010\rangle, |101\rangle, |111\rangle$

r = 110: $|000\rangle, |001\rangle, |110\rangle, |111\rangle$

r = 111: $|000\rangle, |011\rangle, |101\rangle, |110\rangle$

These subspaces are orthogonal except for intersections in 2-dimensional planes. The period can be found by measuring in the computational basis. Repetitions of the measurement will eventually yield sufficiently many distinct values to determine in which subspace out of the seven possibilities the final state lies. In this case ($n = 3$), it is clear by examining the above list that two values distinct from 000 suffice to determine the subspace, and these are just the values $y_i$ for which $y_i \cdot r = 0$. Note that the subspaces correspond to quantum disjunctions. So determining the period of the function by Simon's algorithm amounts to determining which disjunction out of the seven alternative disjunctions is true, i.e., which subspace contains the state, without determining the truth values of the disjuncts.

### 6.2.3 Shor's Algorithm

Shor's factorization algorithm exploits the fact that the two prime factors $p, q$ of a positive integer $N = pq$ can be found by determining the period of a function $f(x) = a^x \bmod N$, for any $a < N$ which is coprime to $N$, i.e., has no common factors with $N$ other than 1. The period $r$ of $f(x)$ depends on $a$ and $N$. Once we know the period, we can factor $N$ if $r$ is even and $a^{r/2} \neq -1 \bmod N$, which will be the case with probability greater than 1/2 if $a$ is chosen randomly. (If not, we choose another value of $a$.) The factors of $N$ are the greatest common factors of $a^{r/2} \pm 1$ and $N$, which can be found in polynomial time by the Euclidean algorithm. (For these number-theoretic results, see [Nielsen and Chuang, 2000, Appendix 4].) So the problem of factorizing a composite integer $N$ that is the product of two primes reduces to the problem of finding the period of a certain periodic function $f : Z_s \to Z_N$, where $Z_n$ is the additive group of integers mod $n$ (rather than $B^n$, the $n$-fold Cartesian product of a Boolean algebra $B$, as in Simon's algorithm). Note that $f(x + r) = f(x)$ if $x + r \leq s$. The function $f$ is periodic if $r$ divides $s$ exactly, otherwise it is almost periodic.

Consider first the general form of the algorithm, as it is usually formulated. We begin by initializing the input register ($s$ qubits) to the state $|0\rangle \in \mathcal{H}^s$ and the output register ($N$ qubits) to the state $|0\rangle \in \mathcal{H}^N$. An $s$-fold Hadamard transformation is applied to the input register, followed by the unitary transformation $U_f$ which implements the function $f(x) = a^x \bmod N$:

$$|0\rangle|0\rangle \quad \stackrel{H}{\longrightarrow} \quad \frac{1}{\sqrt{s}} \sum_{x=0}^{s-1} |x\rangle|0\rangle \tag{253}$$

$$\stackrel{U_f}{\longrightarrow} \quad \frac{1}{\sqrt{s}} \sum_{x=0}^{s-1} |x\rangle|0\rangle \tag{254}$$

$$= \frac{1}{\sqrt{s}} \sum_{x=0}^{s-1} |x\rangle|x + a^x \bmod N\rangle \tag{255}$$

Then we measure the output register in the computational basis[17] and obtain a state of the following form for the input register:

$$\frac{1}{\sqrt{s/r}} \sum_{j=0}^{s/r-1} |x_i + jr\rangle \tag{256}$$

This will be the case if $r$ divides $s$ exactly. The value $x_i$ is the offset, which depends on the outcome $i$ of the measurement of the output register. The sum is taken over the values of $j$ for which $f(x_i + jr) = i$. When $r$ does not divide $s$ exactly, the analysis is a little more complicated. For a discussion, see [Barenco, 1998, p. 164], [Jozsa, 1997b]. Since the state label contains the random offset, a direct measurement of the label yields no information about the period.

[17] As in the discussion of Simon's algorithm, this measurement is purely hypothetical, introduced to simplify the analysis. Only the input register is actually measured.

A discrete Fourier transform for the integers mod $s$ is now applied to the input register, i.e., a unitary transformation:

$$|x\rangle \stackrel{U_{DFT_s}}{\longrightarrow} \frac{1}{\sqrt{s}}\sum_{y=0}^{s-1} e^{2\pi i \frac{xy}{s}}|y\rangle, \text{ for } x \in Z_s \tag{257}$$

This yields the transition:

$$\frac{1}{\sqrt{\frac{s}{r}}}\sum_{j=0}^{\frac{s}{r}-1}|x_i + jr\rangle \stackrel{U_{DFT_s}}{\longrightarrow} \frac{1}{\sqrt{r}}\sum_{k=0}^{r-1} e^{2\pi i \frac{x_i k}{r}}|ks/r\rangle \tag{258}$$

and so shifts the offset into a phase factor and inverts the period as a multiple of $s/r$. A measurement of the input register in the computational basis yields $c = ks/r$. The algorithm is run a number of times until a value of $k$ coprime to $r$ is obtained. Cancelling $c/s$ to lowest terms then yields $k$ and $r$ as $k/r$.

Since we don't know the value of $r$ in advance of applying the algorithm, we do not, of course, recognize when a measurement outcome yields a value of $k$ coprime to $r$. The idea is to run the algorithm, cancel $c/s$ to lowest terms to obtain a candidate value for $r$ and hence a candidate factor of $N$, which can then be tested by division into $N$. Even when we do obtain a value of $k$ coprime to $r$, some values of $a$ will yield a period for which the method fails to yield a factor of $N$, in which case we randomly choose a new value of $a$ and run the algorithm with this value. The point is that all these steps are efficient, i.e., can be performed in polynomial time, and since only a polynomial number of repetitions are required to determine a factor with any given probability $p < 1$, the algorithm is a polynomial-time algorithm, achieving an exponential speed-up over any known classical algorithm.

To see how the algorithm functions geometrically, consider the case $N = 15, a = 7$ and $s = 64$ discussed in [Barenco, 1998, p. 160]. In this case, the function $f(x) = a^x \bmod 15$ is:

$$\begin{aligned} 7^0 \bmod 15 &= 1 \\ 7^1 \bmod 15 &= 7 \\ 7^2 \bmod 15 &= 4 \\ 7^3 \bmod 15 &= 13 \\ 7^4 \bmod 15 &= 1 \\ &\vdots \end{aligned}$$

and the period is evidently $r = 4$.[18] After the application of the unitary transformation $U_f = a^x \bmod N$, the state of the two registers is:

$$\begin{aligned} &\tfrac{1}{8}(|0\rangle|1\rangle + |1\rangle|7\rangle + |2\rangle|4\rangle + |3\rangle|13\rangle \\ &+ |4\rangle|1\rangle + |5\rangle|7\rangle + |6\rangle|4\rangle + |7\rangle|13\rangle \end{aligned}$$

[18]The factors 3 and 5 of 15 are derived as the greatest common factors of $a^{r/2} - 1 = 48$ and 15 and $a^{r/2} + 1 = 50$ and 15, respectively.

$$\vdots$$
$$+ |60\rangle|1\rangle + |61\rangle|7\rangle + |62\rangle|4\rangle + |63\rangle|13\rangle) \quad (259)$$

This state can be expressed as:

$$\begin{aligned} &\tfrac{1}{4}(|0\rangle + |4\rangle + |8\rangle + \ldots + |60\rangle)|1\rangle \\ &+\tfrac{1}{4}(|1\rangle + |5\rangle + |9\rangle + \ldots + |61\rangle)|7\rangle \\ &+\tfrac{1}{4}(|2\rangle + |6\rangle + |10\rangle + \ldots + |62\rangle)|4\rangle \\ &+\tfrac{1}{4}(|3\rangle + |7\rangle + |11\rangle + \ldots + |63\rangle)|13\rangle) \end{aligned} \quad (260)$$

If we measure the output register, we obtain (equiprobably) one of four states for the input register, depending on the outcome of the measurement: 1, 7, 4, or 13:

$$\tfrac{1}{4}(|0\rangle + |4\rangle + |8\rangle + \ldots + |60\rangle) \quad (261)$$
$$\tfrac{1}{4}(|1\rangle + |5\rangle + |9\rangle + \ldots + |61\rangle) \quad (262)$$
$$\tfrac{1}{4}(|2\rangle + |6\rangle + |10\rangle + \ldots + |62\rangle) \quad (263)$$
$$\tfrac{1}{4}(|3\rangle + |7\rangle + |11\rangle + \ldots + |63\rangle) \quad (264)$$

These are the states (256) for values of the offset 0, 1, 2, 3. Application of the discrete Fourier transform yields:

$x_1 = 0 : \ \tfrac{1}{2}(|0\rangle + |16\rangle + |32\rangle + |48\rangle)$

$x_7 = 1 : \ \tfrac{1}{2}(|0\rangle + i|16\rangle - |32\rangle - i|48\rangle)$

$x_4 = 2 : \ \tfrac{1}{2}(|0\rangle - |16\rangle + |32\rangle - |48\rangle)$

$x_{13} = 3 : \ \tfrac{1}{2}(|0\rangle - i|16\rangle - |32\rangle + i|48\rangle)$

which are the states in (258). So for the period $r = 4$, the state of the input register ends up in the 4-dimensional subspace spanned by the vectors $|0\rangle, |16\rangle, |32\rangle, |48\rangle$.

Now consider all possible even periods $r$ for which $f(x) = a^x \bmod 15$, where $a$ is coprime to $15$. The other possible values of $a$ are 2, 4, 8, 11, 13, 14 and the corresponding periods turn out to be 4, 2, 4, 2, 4, 2. So we need only consider $r = 2$. [19]

For $r = 2$, if we measure the output register, we will obtain (equiprobably) one of two states for the input register, depending on the outcome of the measurement (say, $a$ or $b$):

$$|0\rangle + |2\rangle + |4\rangle + \ldots + |62\rangle \quad (265)$$
$$|1\rangle + |3\rangle + |5\rangle + \ldots + |63\rangle \quad (266)$$

After the discrete Fourier transform, these states are transformed to:

$x_a = 0 : \ |0\rangle + |32\rangle$

[19] Every value of $a$ except $a = 14$ yields the correct factors for 15. For $a = 14$, the method fails: $r = 2$, so $a^{\frac{r}{2}} = -1 \bmod 15$.

$x_b = 1 : \; |0\rangle - |32\rangle$

In this case, the 2-dimensional subspace $\mathcal{V}_{r=2}$ spanned by $|0\rangle, |32\rangle$ for $r = 2$ is included in the 4-dimensional subspace $\mathcal{V}_{r=4}$ for $r = 4$. A measurement can distinguish $r = 4$ from $r = 2$ reliably, i.e., whether the final state of the input register is in $\mathcal{V}_{r=4}$ or $\mathcal{V}_{r=2}$, only if the final state is in $\mathcal{V}_{r=4} - \mathcal{V}_{r=2}$, the part of $\mathcal{V}_{r=4}$ orthogonal to $\mathcal{V}_{r=2}$. What happens if the final state ends up in $\mathcal{V}_{r=2}$?

Shor's algorithm works as a randomized algorithm. As mentioned above, it produces a candidate value for the period $r$ and hence a candidate factor of $N$, which can be tested (in polynomial time) by division into $N$. A measurement of the input register in the computational basis yields an outcome $c = ks/r$. The value of $k$ is chosen equiprobably by the measurement of the output register. The procedure is to repeat the algorithm until the outcome yields a value of $k$ coprime to $r$, in which case canceling $c/s$ to lowest terms yields $k$ and $r$ as $k/r$.

For example, suppose we choose $a = 7$, in which case (unknown to us) $r = 4$. The values of $k$ coprime to $r$ are $k = 1$ and $k = 3$ (this is also unknown to us, because $k$ depends on the value of $r$). Then $c/s$ cancelled to lowest terms is $1/4$ and $3/4$, respectively, both of which yield the correct period. From the geometrical perspective, these values of $k$ correspond to finding the state after measurement in the computational basis to be $|16\rangle$ or $|48\rangle$, both of which do distinguish $\mathcal{V}_{r=4}$ from $\mathcal{V}_{r=2}$.

Suppose we choose a value of $a$ with period $r = 2$ and find the value $c = 32$. The only value of $k$ coprime to $r$ is $k = 1$. Then $c/s$ cancelled to lowest terms is $1/2$, which yields the correct period, and hence the correct factors of $N$. But $c = 32$ could also be obtained for $a = 7$, $r = 4$, and $k = 2$, which does not yield the correct period, and hence does not yield the correct factors of $N$. Putting it geometrically: the value $k = 1$ for $r = 2$ corresponds to the same state, $|32\rangle$, as the value $k = 2$ for $r = 4$. Once we obtain the candidate period $r = 2$ (by cancelling $c/s = 32/64$ to lowest terms), we calculate the factors of $N$ as the greatest common factors of $a \pm 1$ and $N$ and test these by division into $N$. If $a = 7$, these calculated factors will be incorrect. If $a = 2$, say, the factors calculated in this way will be correct.

We see that, with the added information provided by the outcome of a test division of a candidate factor into $N$, Shor's randomized algorithm again amounts to determining which disjunction among alternative disjunctions is true, i.e., which subspace contains the state, without determining the truth values of the disjuncts.

## 6.3 Where Does the Speed-Up Come From?

What, precisely, is the feature of a quantum computer responsible for the phenomenal efficiency over a classical computer? In the case of Simon's algorithm, the speed-up is exponential over any classical algorithm; in the case of Shor's algorithm, the speed-up is exponential over any known classical algorithm.

Steane [1998] remarks:

> The period finding algorithm appears at first sight like a conjuring trick: it is not quite clear how the quantum computer managed to produce the period like a rabbit out of a hat. . . . I would say that the most important features are contained in [$|\psi\rangle = \frac{1}{s}\sum_{x=0}^{s-1} |x\rangle |f(x)\rangle$]. They are not only the

> *quantum parallelism* already mentioned, but also *quantum entanglement*, and, finally, quantum interference. Each value of $f(x)$ retains a link with the value of $x$ which produced it, through entanglement of the $x$ and $y$ registers in [$|\psi\rangle$]. The 'magic' happens when a measurement of the $y$ register produces the special state [$\frac{1}{s/r}\sum_{j=0}^{s/r-1}|x_i + jr\rangle$] in the $x$-register, and it is quantum entanglement which permits this (see also [Jozsa, 1997a]). The final Fourier transform can be regarded as an interference between the various superposed states in the $x$-register (compare with the action of a diffraction grating).
>
> Interference effects can be used for computational purposes with classical light fields, or water waves for that matter, so interference is not in itself the essentially quantum feature. Rather, the exponentially large number of interfering states, and the entanglement, are features which do not arise in classical systems.

Jozsa points out [1997a] that the state space (phase space) of a composite classical system is the Cartesian product of the state spaces of its subsystems, while the state space of a composite quantum system is the tensor product of the state spaces of its subsystems. For $n$ qubits, the quantum state space has $2^n$ dimensions. So the information required to represent a general state increases exponentially with $n$: even if we restrict the specification of the amplitudes to numbers of finite precision, a superposition will in general have $\mathcal{O}(2^n)$ components. For a classical composite system of $n$ two-level subsystems, the number of possible states grows exponentially with $n$, but the information required to represent a general state is just $n$ times the information required to represent a single two-level system, i.e., the information grows only linearly with $n$ because the state of a composite system is just a product state.

More formally, Jozsa and Linden [2002] have shown that a quantum algorithm operating on pure states can achieve an exponential speed-up over classical algorithms only if the quantum algorithm involves multi-partite entanglement that increases unboundedly with the input size. Similarly, Vidal [2003] has shown that a classical computer can simulate the evolution of a pure state of $n$ qubits with computational resources that grow linearly with $n$ and exponentially in multi-partite entanglement.

The essential feature of the quantum computations discussed above in §6.2 is the selection of a disjunction, representing a global property of a function, among alternative possible disjunctions without computing the truth values of the disjuncts, which is redundant information in a quantum computation but essential information classically. Note that a quantum disjunction is represented by a subspace of entangled states in the tensor product Hilbert space of the input and output registers. This is analogous to the procedure involved in the key observation underlying the proof of the quantum bit commitment theorem discussed in §5.2.2. The series of operations described by equations (167)–(172), in which the channel particle is entangled with ancilla systems and the ancillas are subsequently measured, effectively constitute a quantum computaton.

The first stage of a quantum algorithm involves the creation of a state in which every input value to the function is correlated with a corresponding output value. This is referred to as 'quantum parallelism' and is sometimes cited as the source of the speed-up in a quantum computation. The idea is that a quantum computation is something

like a massively parallel classical computation, for all possible values of a function. This appears to be Deutsch's view, with the parallel computations taking place in parallel universes. For a critique, see [Steane, 2003], who defends a view similar to that presented here. Of course, all these different values are inaccessible: a measurement in the computational basis will only yield (randomly) one correlated input-output pair. Further processing is required, including the final discrete Fourier transform for the three algorithms discussed in §6.2. It would be incorrect to attribute the efficiency of these quantum algorithms to the interference in the input register produced by the Fourier transform. The role of the Fourier transform is simply to allow a measurement in the computational basis to reveal which subspace representing the target disjunction contains the state.

One might wonder, then, why the discrete Fourier transform is even necessary. We could, of course, simply perform an equivalent measurement in a different basis. But note that a computation would have to be performed to determine this basis. This raises the question of precisely how to assess the speed-up of a quantum algorithm relative to a rival classical algorithm. What are the relevant computational steps to be counted in making this assessment for a quantum computation? Since any sequence of unitary transformations is equivalent to a single unitary transformation, and a unitary transformation followed by a measurement in a certain basis is equivalent to simply performing a measurement in a different basis, any quantum computation can always be reduced to just one step: a measurement in a particular basis!

Of course, this observation is hardly illuminating, since a computation at least as difficult as the original computation would have to be performed to determine the required basis, but it does indicate that some convention is required about what steps to count in a quantum computation. The accepted convention is to require the unitary transformations in a quantum computation to be constructed from elementary quantum gates that form a universal set (e.g., the CNOT gate, the Hadamard gate, the phase gate, and the $\pi/8$ gate discussed in §6.1) and to count each such gate as one step. In addition, all measurements are required to be performed in the computational basis, and these are counted as additional steps. The final discrete Fourier transforms in the Deutsch-Jozsa algorithm, Simon's algorithm, and Shor's algorithm are indispensable in transforming the state so that the algorithms can be completed by measurements in the computational basis, and it is an important feature of these algorithms that the Fourier transform can be implemented efficiently with elementary unitary gates. To claim that a quantum algorithm is exponentially faster than a classical algorithm is to claim that the number of steps counted in this way for the quantum algorithm is a polynomial function of the size of the input (the number of qubits required to store the input), while the classical algorithm involves a number of steps that increases exponentially with the size of the input (the number of bits required to store the input).

# 7 Quantum Foundations from the Perspective of Quantum Information

Does the extension of the classical theory of information to quantum states shed new light on the foundational problems of quantum mechanics underlying the Bohr-Einstein debate mentioned in §1, in particular the measurement problem? Researchers in the area of quantum information and quantum computation often suggest a positive answer to this question, with a promissory note for how the story is supposed to go. More fully worked-out (generally, rather different) positive responses have been proposed by various authors, notably Fuchs [2001b, 2002a, 2002b, 2001a] and Brukner and Zeilinger [2001, 20002]. For a very thorough analysis and critique of the Brukner-Zeilinger position, see [Timpson, 2004]. See also Hall [2000] and the response by Brukner and Zeilinger [2000]. Here I shall limit my discussion to the significance of a characterization theorem for quantum mechanics in terms of information-theoretic constraints by Clifton, Bub, and Halvorson (CBH) [2003].

## 7.1 The CBH Characterization Theorem

CBH showed that one can derive the basic kinematic features of a quantum description of physical systems from three fundamental information-theoretic constraints:

- the impossibility of superluminal information transfer between two physical systems by performing measurements on one of them,
- the impossibility of perfectly broadcasting the information contained in an unknown physical state (which, for pure states, amounts to 'no cloning'),
- the impossibility of communicating information so as to implement a bit commitment protocol with unconditional security (so that cheating is in principle excluded by the theory).

More precisely, CBH formulate these information-theoretic constraints in the general framework of $C^*$-algebras, which allows a mathematically abstract characterization of a physical theory that includes, as special cases, all classical mechanical theories of both wave and particle varieties, and all variations on quantum theory, including quantum field theories (plus any hybrids of these theories, such as theories with superselection rules). Within this framework, CBH show that the three information-theoretic constraints jointly entail three physical conditions that they take as definitive of what it means to be a quantum theory in the most general sense. Specifically, the information-theoretic constraints entail that:

- the algebras of observables pertaining to distinct physical systems commute (a condition usually called microcausality or, to use Summers' term [Summers, 1990], *kinematic independence*),
- any individual system's algebra of observables is *noncommutative*,
- the physical world is *nonlocal*, in that spacelike separated systems can occupy entangled states that persist as the systems separate.

CBH also partly demonstrated the converse derivation, leaving open a question concerning nonlocality and bit commitment. This remaining issue was later resolved by Hans Halvorson [2004], so the CBH theorem is a characterization theorem for quantum theory in terms of the three information-theoretic constraints.

Note that the $C^*$-algebraic framework is not restricted to the standard quantum mechanics of a system represented on a single Hilbert space with a unitary dynamics, but is general enough to cover cases of systems with an infinite number of degrees of freedom that arise in quantum field theory and the thermodynamic limit of quantum statistical mechanics (in which the number of microsystems and the volume they occupy goes to infinity, while the density defined by their ratio remains constant). The $C^*$-algebraic framework has even been applied to quantum field theory on curved spacetime and so is applicable to the quantum theoretical description of exotic phenomena such as Hawking radiation (black hole evaporation); see [Wald, 1984]. The Stone-von Neumann theorem, which guarantees the existence of a unique representation (up to unitary equivalence) of the canonical commutation relations for systems with a finite number of degrees of freedom, breaks down for such cases, and there will be many unitarily inequivalent representations of the canonical commutation relations.

One could, of course, consider weaker mathematical structures, but it seems that the $C^*$-algebraic machinery suffices for all physical theories that have been found to be empirically successful to date, including phase space theories and Hilbert space theories [Landsman, 1998], and theories based on a manifold [Connes, 1994]. For further discussion of this point, see Halvorson and Bub [2005]. See also Halvorson (this vol., chap. 8), Emch (this vol., ch. 8), and Landsman (this vol., ch. 5).

A $C^*$-algebra is essentially an abstract generalization of the structure of the algebra of operators on a Hilbert space. Technically, a (unital) $C^*$-algebra is a Banach $^*$-algebra over the complex numbers containing the identity, where the involution operation $^*$ and the norm are related by $\|A^*A\| = \|A\|^2$. So the algebra $\mathfrak{B}(\mathcal{H})$ of all bounded operators on a Hilbert space $\mathcal{H}$ is a $C^*$-algebra, with $^*$ the adjoint operation and $\|\cdot\|$ the standard operator norm.

In standard quantum theory, as discussed in §3.1.1, a state on $\mathfrak{B}(\mathcal{H})$ is defined by a density operator $\rho$ on $\mathcal{H}$ in terms of an expectation-valued functional $\tilde{\rho}(A) = \mathrm{Tr}(\rho A)$ for all observables represented by self-adjoint operators $A$ in $\mathfrak{B}(\mathcal{H})$. This definition of $\tilde{\rho}(A)$ in terms of $\rho$ yields a positive normalized linear functional. So a state on a $C^*$-algebra $\mathfrak{C}$ is defined, quite generally, as any positive normalized linear functional $\tilde{\rho} : \mathfrak{C} \to \mathbb{C}$ on the algebra. Pure states can be defined by the condition that if $\tilde{\rho} = \lambda\tilde{\rho}_1 + (1-\lambda)\tilde{\rho}_2$ with $\lambda \in (0,1)$, then $\tilde{\rho} = \tilde{\rho}_1 = \tilde{\rho}_2$; other states are mixed. In the following, we drop the '~' in $\tilde{\rho}$, but note that a $C^*$-algebraic state $\rho$ is a positive linear functional on $\mathfrak{C}$, while the density operator of standard quantum mechanics is an element of $\mathfrak{C} = \mathfrak{B}(\mathcal{H})$.

By Gleason's theorem [Gleason, 1957], every $C^*$-algebraic state in this sense on a $C^*$-algebra $\mathfrak{C} = \mathfrak{B}(\mathcal{H})$ is given by a density operator on $\mathfrak{B}(\mathcal{H})$. However, because countable additivity is not presupposed by the $C^*$-algebraic notion of state (and, therefore, Gleason's theorem does not apply in general), there can be pure states of $\mathfrak{B}(\mathcal{H})$ that are not representable by vectors in $\mathcal{H}$. In fact, if $A$ is any self-adjoint element of a $C^*$-algebra $\mathfrak{A}$, and $a \in \mathrm{sp}(A)$, then there always exists a pure state $\rho$ of $\mathfrak{A}$ that assigns a *dispersion-free* value of $a$ to $A$ [Kadison and Ringrose, 1997, Ex. 4.6.31]. Since this

is true even when we consider a point in the continuous spectrum of a self-adjoint operator $A$ acting on a Hilbert space, *without* any corresponding eigenvector, it follows that there *are* pure states of $\mathfrak{B}(\mathcal{H})$ in the $C^*$-algebraic sense that cannot be vector states (nor, in fact, representable by any density operator $\mathcal{H}$).

As we saw in §3.1.3, the general evolution of a quantum system resulting from a combination of unitary interactions and selective or nonselective measurements can be described by a quantum operation, i.e., a completely positive linear map. Accordingly, a completely positive linear map $T : \mathfrak{C} \to \mathfrak{C}$, where $0 \leq T(I) \leq I$ is taken as describing the general evolution of a system represented by a $C^*$-algebra of observables. The map or operation $T$ is called selective if $T(I) < I$ and nonselective if $T(I) = I$. Recall that a yes-no measurement of some idempotent observable represented by a projection operator $P$ is an example of a selective operation. Here, $T(A) = PAP$ for all $A$ in the $C^*$-algebra $\mathfrak{C}$, and $\rho^T$, the transformed ('collapsed') state, is the final state obtained after measuring $P$ in the state $\rho$ and ignoring all elements of the ensemble that do not yield the eigenvalue 1 of $P$ (so $\rho^T(A) = \rho(T(A))/\rho(T(I))$ when $\rho(T(I)) \neq 0$, and $\rho^T = 0$ otherwise). The time evolution in the Heisenberg picture induced by a unitary operator $U \in \mathfrak{C}$ is an example of a nonselective operation. Here, $T(A) = UAU^{-1}$. Similarly, the measurement of an observable $O$ with spectral measure $\{P_i\}$, without selecting a particular outcome, is an example of a nonselective operation, with $T(A) = \sum_{i=1}^{n} P_i A P_i$. As in the standard quantum theory of a system with a finite-dimensional Hilbert space (cf. Eq. (67) of §3.1.3), any completely positive linear map can be regarded as the restriction to a local system of a unitary map on a larger system.

A representation of a $C^*$-algebra $\mathfrak{C}$ is any mapping $\pi : \mathfrak{C} \to \mathfrak{B}(\mathcal{H})$ that preserves the linear, product, and $^*$ structure of $\mathfrak{C}$. The representation is faithful if $\pi$ is one-to-one, in which case $\pi(\mathfrak{C})$ is an isomorphic copy of $\mathfrak{C}$. The Gelfand-Naimark theorem says that every abstract $C^*$-algebra has a concrete faithful representation as a norm-closed $^*$-subalgebra of $\mathfrak{B}(\mathcal{H})$, for some appropriate Hilbert space $\mathcal{H}$. As indicated above, in the case of systems with an infinite number of degrees of freedom (e.g., quantum field theory), there are inequivalent representations of the $C^*$-algebra of observables defined by the commutation relations.

Every classical phase space theory defines a commutative $C^*$-algebra. For example, the observables of a classical system of $n$ particles—the real-valued continuous functions on the phase space $\mathbb{R}^{6n}$—can be represented as the self-adjoint elements of the $C^*$-algebra $\mathfrak{B}(\mathbb{R}^{6n})$ of all continuous complex-valued functions $f$ on $\mathbb{R}^{6n}$. The phase space $\mathbb{R}^{6n}$ is locally compact and can be made compact by adding just one point 'at infinity,' or we can simply consider a bounded (and thus compact) subset of $\mathbb{R}^{6n}$. The statistical states of the system are given by probability measures $\mu$ on $\mathbb{R}^{6n}$, and pure states, corresponding to maximally complete information about the particles, are given by the individual points of $\mathbb{R}^{6n}$. The system's state $\rho$ in the $C^*$-algebraic sense is the expectation functional corresponding to $\mu$, defined by $\rho(f) = \int_{\mathbb{R}^{6n}} f \mathrm{d}\mu$. Conversely [Kadison and Ringrose, 1997, Thm. 4.4.3], every *commutative* $C^*$-algebra $\mathfrak{C}$ is isomorphic to the set $C(X)$ of all continuous complex-valued functions on a locally compact Hausdorff space $X$ defined by the pure states of $\mathfrak{C}$. If $\mathfrak{C}$ has a multiplicative identity, the 'phase space' $X$ is compact. In this 'function representation' of $C$, the isomorphism maps an element $C \in \mathfrak{C}$ to the function $\hat{C}$ (the Gelfand transformation

of $C$) whose value at any $\rho$ is just the (dispersion-free) value that $\rho$ assigns to $C$. So 'behind' every abstract commutative $C^*$-algebra there is a classical phase space theory defined by its function representation on the phase space $X$. This representation theorem (and its converse) justifies characterizing a $C^*$-algebraic theory as classical just in case its algebra is commutative.

As we saw above, CBH identify quantum theories with a certain subclass of noncommutative $C^*$-algebras, where the condition of kinematic independence is satisfied by the algebras of observables of distinct systems and the states of spacelike separated systems are characterized by the sort of nonlocality associated with entanglement.

To clarify the rationale for this characterization and the significance of the information-theoretic constraints, consider a composite quantum system $AB$, consisting of two subsystems, $A$ and $B$. For simplicity, assume the systems are indistinguishable, so their $C^*$-algebras $\mathfrak{A}$ and $\mathfrak{B}$ are isomorphic. The observables of the component systems $A$ and $B$ are represented by the self-adjoint elements of $\mathfrak{A}$ and $\mathfrak{B}$, respectively. Let $\mathfrak{A}\vee\mathfrak{B}$ denote the $C^*$-algebra generated by $\mathfrak{A}$ and $\mathfrak{B}$. The physical states of $A$, $B$, and $AB$, are given by positive normalized linear functionals on their respective algebras that encode the expectation values of all observables. To capture the idea that $A$ and $B$ are *physically distinct* systems, CBH make the assumption that any state of $\mathfrak{A}$ is compatible with any state of $\mathfrak{B}$, i.e., for any state $\rho_A$ of $\mathfrak{A}$ and $\rho_B$ of $\mathfrak{B}$, there is a state $\rho$ of $\mathfrak{A} \vee \mathfrak{B}$ such that $\rho|_{\mathfrak{A}} = \rho_A$ and $\rho|_{\mathfrak{B}} = \rho_B$.

The sense of the 'no superluminal information transfer via measurement' constraint is that when Alice and Bob, say, perform local measurements, Alice's measurements can have no influence on the statistics for the outcomes of Bob's measurements, and conversely. That is, merely performing a local measurement cannot, in and of itself, convey any information to a physically distinct system, so that everything 'looks the same' to that system after the measurement operation as before, in terms of the expectation values for the outcomes of measurements. CBH show [2003, Thm. 1] that it follows from this constraint that A and B are *kinematically independent* systems if they are physically distinct in the above sense, i.e., every element of $\mathfrak{A}$ commutes pairwise with every element of $\mathfrak{B}$. (More precisely, an operation $T$ on $\mathfrak{A} \vee \mathfrak{B}$ conveys no information to Bob just in case $(T^*\rho)|_{\mathfrak{B}} = \rho|_{\mathfrak{B}}$ for all states $\rho$ of $\mathfrak{B}$, where $T^*$ is the map on the states, i.e., the positive linear functionals on $\mathfrak{A} \vee \mathfrak{B}$, induced by T. Clearly, the kinematic independence of $\mathfrak{A}$ and $\mathfrak{B}$ entails that Alice's local measurement operations cannot convey any information to Bob, i.e., $T(B) = \sum_{i=1}^{n} E_i^{1/2} B E_i^{1/2} = B$ for $B \in \mathfrak{B}$ if $T$ is implemented by a POVM in $\mathfrak{A}$. CBH prove that if Alice cannot convey any information to Bob by performing local measurement operations, then $\mathfrak{A}$ and $\mathfrak{B}$ are kinematically independent.)

The 'no broadcasting' condition now ensures that the individual algebras $\mathfrak{A}$ and $\mathfrak{B}$ are noncommutative. Recall that for pure states, broadcasting reduces to cloning, and that in elementary quantum mechanics, neither cloning nor broadcasting is possible in general (see section 3.2). CBH show that broadcasting and cloning are always possible for classical systems, i.e., in the commutative case there *is* a universal broadcasting map that clones any pair of input pure states and broadcasts any pair of input mixed states [Clifton *et al.*, 2003, Thm. 2]. Conversely, they show that if any two states can be (perfectly) broadcast, then any two pure states can be cloned; and if two pure states

of a $C^*$-algebra can be cloned, then they must be orthogonal. So, if any two states can be broadcast, then all pure states are orthogonal, from which it follows that the algebra is commutative.

The quantum mechanical phenomenon of interference is the physical manifestation of the noncommutativity of quantum observables or, equivalently, the superposition of quantum states. So the impossibility of perfectly broadcasting the information contained in an unknown physical state, or of cloning or copying the information in an unknown pure state, is the information-theoretic counterpart of interference.

Now, if $\mathfrak{A}$ and $\mathfrak{B}$ are noncommutative and mutually commuting, it can be shown that there are nonlocal entangled states on the $C^*$-algebra $\mathfrak{A} \vee \mathfrak{B}$ they generate (see [Landau, 1987, Summers, 1990, Bacciagaluppi, 1994], and—more relevantly here, in terms of a specification of the range of entangled states that can be guaranteed to exist—[Halvorson, 2004]). So it seems that entanglement—what Schrödinger [1935, p. 555] identified as '*the* characteristic trait of quantum mechanics, the one that enforces its entire departure from classical lines of thought,' as we saw in §4.1— follows automatically in any theory with a noncommutative algebra of observables. That is, it seems that once we assume 'no superluminal information transfer via measurement,' and 'no broadcasting,' the class of allowable physical theories is restricted to those theories in which physical systems manifest both interference *and* nonlocal entanglement. But in terms of physical interpretation this conclusion is a bit too quick, since the derivation of entangled states depends on formal properties of the $C^*$-algebraic machinery. Moreover, we have no assurance that two systems in an entangled state will maintain their entanglement indefinitely as they separate in space, which is the case for quantum entanglement. But this is precisely what is required by the cheating strategy that thwarts secure bit commitment, since Alice will have to keep one system of such a pair and send the other system to Bob, whose degree of spatial separation from Alice is irrelevant, in principle, to the implementation of the protocol. In an information-theoretic characterization of quantum theory, the fact that entangled states of composite systems can be instantiated, and instantiated nonlocally so that the entanglement of composite system is maintained as the subsystems separate in space, should be shown to follow from some information-theoretic principle. The role of the 'no bit commitment' constraint is to guarantee the persistence of entanglement over distance, i.e., the existence of a certain class of nonlocal entangled states—hence it gives us nonlocality, not merely 'holism.'

As shown in §5.2, unconditionally secure quantum bit commitment is impossible because a generalized version of the EPR cheating strategy can always be applied by introducing additional ancilla particles and enlarging the Hilbert space in a suitable way. That is, for a quantum mechanical system consisting of two (separated) subsystems represented by the $C^*$-algebra $\mathfrak{B}(\mathcal{H}_1) \otimes \mathfrak{B}(\mathcal{H}_2)$, any mixture of states on $\mathfrak{B}(\mathcal{H}_2)$ can be generated from a distance by performing an appropriate generalized measurement on the system represented by $\mathfrak{B}(\mathcal{H}_1)$, for an appropriate entangled state of the composite system. This is what Schrödinger called 'remote steering' and found so physically counterintuitive that he speculated [1936, p. 451] (wrongly, as it turned out) that experimental evidence would eventually show that this was simply an artefact of the theory, not instantiated in our world. He suggested that an entangled state of a composite system would almost instantaneously decay to a mixture as the compo-

nent systems separated.[20] There would still be correlations between the states of the component systems, but remote steering would no longer be possible.

> It seems worth noticing that the [EPR] paradox could be avoided by a very simple assumption, namely if the situation after separating were described by [the entangled state $\Psi(x, y) = \sum_k a_k g_k(x) f_k(y)$], but with the additional statement that the knowledge of the *phase relations* between the complex constants $a_k$ has been entirely lost in consequence of the process of separation. This would mean that not only the parts, but the whole system, would be in the situation of a mixture, not of a pure state. It would not preclude the possibility of determining the state of the first system by *suitable* measurements in the second one or *vice versa*. But it would utterly eliminate the experimenters influence on the state of that system which he does not touch.

Schrödinger regarded the phenomenon of interference associated with noncommutativity in quantum mechanics as unproblematic, because he saw this as reflecting the fact that particles are wavelike. But he did not believe that we live in a world in which physical systems can exist nonlocally in entangled states, because such states would allow Alice to steer Bob's system into any mixture of pure states compatible with Bob's reduced density operator and he did not expect that experiments would bear this out. Of course, it was an experimental question in 1935 whether Schrödinger's conjecture was correct or not. We now know that the conjecture is false. A wealth of experimental evidence, including the confirmed violations of Bell's inequality [Aspect *et al.*, 1981, Aspect *et al.*, 1982] and the confirmations of quantum teleportation [Bouwmeester *et al.*, 1997, Boschi *et al.*, 1998, Furasawa *et al.*, 1998, Nielsen *et al.*, 1998], testify to this. The relevance of Schrödinger's conjecture here is this: it raises the possibility of a quantum-like world in which there is interference but no nonlocal entanglement. Can we exclude this possibility on information-theoretic grounds?

Now although unconditionally secure bit commitment is no less impossible for classical systems, in which the algebras of observables are commutative, than for quantum systems, the insecurity of a bit commitment protocol in a noncommutative setting depends on considerations entirely different from those in a classical commutative setting. As we saw in §5.2, the security of a classical bit commitment protocol is a matter of computational complexity and cannot be unconditional.

By contrast, if, as Schrödinger speculated, we lived in a world in which the algebras of observables are noncommutative but composite physical systems cannot exist in nonlocal entangled states, if Alice sends Bob one of two mixtures associated with the same density operator to establish her commitment, then she is, in effect, sending Bob evidence for the truth of an exclusive disjunction that is not based on the selection of a particular disjunct. (Bob's reduced density operator is associated ambiguously with both mixtures, and hence with the truth of the exclusive disjunction: '0 or 1'.) Noncommutativity allows the possibility of different mixtures associated with the same density operator. What thwarts the possibility of using the ambiguity of mixtures in this way to implement an unconditionally secure bit commitment protocol is the existence

[20] A similar possibility was raised and rejected by Furry [1936].

of nonlocal entangled states between Alice and Bob. This allows Alice to cheat by preparing a suitable entangled state instead of one of the mixtures, where the reduced density operator for Bob is the same as that of the mixture. Alice is then able to steer Bob's systems remotely into either of the two mixtures associated with the alternative commitments at will.

So what *would* allow unconditionally secure bit commitment in a noncommutative theory is the absence of physically occupied nonlocal entangled states, or the spontaneous destruction of entanglement as systems separate. One can therefore take Schrödinger's remarks as relevant to the question of whether or not secure bit commitment is possible in our world. In effect, Schrödinger raised the possibility that we live in a quantum-like world in which unconditionally secure bit commitment is possible! It follows that the impossibility of unconditionally secure bit commitment entails that, for any mixed state that Alice and Bob can prepare by following some (bit commitment) protocol, there is a corresponding nonlocal entangled state that can be physically occupied by Alice's and Bob's particles and persists indefinitely as the particles move apart.

To sum up: the content of the CBH theorem is that a quantum theory—a $C^*$-algebraic theory whose observables and states satisfy conditions of kinematic independence, noncommutativity, and nonlocality—can be characterized by three information-theoretic constraints: no superluminal communication of information via measurement, no (perfect) broadcasting, and no (unconditionally secure) bit commitment.

## 7.2 Quantum Mechanics as a Theory of Information

Consider Einstein's view,[21] mentioned in §1, that quantum mechanics is incomplete. Essentially, Einstein based his argument for this claim on the demand that a complete physical theory should satisfy certain principles of realism (essentially, a locality principle and a separability principle), which amounts to the demand that statistical correlations between spatially separated systems should have a common causal explanation in terms of causal factors obtaining at the common origin of the systems. Roughly thirty years after the publication of the Einstein-Podolsky-Rosen paper [1935], John Bell [1964] showed that the statistical correlations of the entangled Einstein-Podolsky-Rosen state for spatially separated particles are inconsistent with any explanation in terms of a classical probability distribution over common causal factors originating at the source of the particles before they separate. But the fact that quantum mechanics allows the possibility of correlations that are not reducible to common causes is a *virtue* of the theory. It is precisely the nonclassical correlations of entangled states that underlie the possibility of an exponential speed-up of quantum computation over classical computation, the possibility of unconditionally secure key distribution but the impossibility of unconditionally secure quantum bit commitment, and phenomena such as quantum teleportation and other nonclassical entanglement-assisted communication protocols.

While Einstein's argument for incompleteness fails, there is another sense, also

[21] The following discussion is adapted from [Bub, 2004] and [Bub, 2005], but the argument here is developed somewhat differently.

associated with entangled states, in which quantum mechanics might be said to be incomplete. In a typical (idealized) quantum mechanical measurement interaction, say an interaction in which the two possible values, $0$ and $1$, of an observable of a qubit in a certain quantum state become correlated with the two possible positions of a macroscopic pointer observable, $p_0$ and $p_1$, the final state is an entangled state, a linear superposition of the states $|0\rangle|p_0\rangle$ and $|1\rangle|p_1\rangle$ with coefficients derived from the initial quantum state of the qubit. To dramatize the problem, Schrödinger [1936] considered the case where $|p_0\rangle$ and $|p_1\rangle$ represent the states of a cat being alive and a cat being dead in a closed box, which is only opened by the observer some time after the measurement interaction. On the standard way of relating the quantum state of a system to what propositions about the system are determinately (definitely) true or false, and what propositions have no determinate truth value, some correlational proposition about the composite system (microsystem $+$ cat) is true in this entangled state, but the propositions asserting that the cat is alive (and the value of the qubit observable is $0$), or that the cat is dead (and the value of the qubit observable is $1$), are assigned no determinate truth value. Moreover, if we assume that the quantum propositions form an algebraic structure isomorphic to the structure of subspaces of the Hilbert space of the composite system—the representational space for quantum states and observables—then it is easy to derive a formal contradiction from the assumption that the correlational proposition corresponding to the entangled state is true, and that the cat is either definitely alive or definitely dead.[22] Schrödinger thought that it was absurd to suppose that quantum mechanics requires us to say that the cat in such a situation (a macrosystem) is neither alive nor dead (does not have a determinate macroproperty of this sort) until an observer opens the box and looks, in which case the entangled state 'collapses' non-linearly and stochastically, with probabilities given by the initial quantum state of the microsystem, onto a product of terms representing a definite state of the cat and a definite state of the microsystem. Einstein [1967, p. 39] concurred and remarked in a letter to Schrödinger: 'If that were so then physics could only claim the interest of shopkeepers and engineers; the whole thing would be a wretched bungle.'

This is the standard 'measurement problem' of quantum mechanics. Admittedly, the formulation of the problem is highly idealized, but the fundamental problem arises from the way in which quantum mechanics represents correlations via entangled states and does not disappear entirely in less idealized formulations (even though the problem is somewhat altered by considering the macroscopic nature of the instrument, and the rôle of the environment). (See Dickson, this vol., ch. 4, and [Bub, 1997] for further discussion.) I shall refer to this problem—the Schrödinger incompleteness of the theory—as Schrödinger's problem. It is a *problem about truth* (or the instantiation of properties), as opposed to a distinct *problem about probabilities*.

Before formulating the probability problem, consider what was involved in the transition from classical to quantum mechanics. Quantum mechanics first appeared as Heisenberg's matrix mechanics in 1925, following the 'old quantum theory,' a patchwork of *ad hoc* modifications of classical mechanics to accommodate Planck's quantum postulate. Essentially, Heisenberg modified the kinematics of classical mechanics

[22]This also follows from the Bub-Clifton theorem discussed below. The sublattice of determinate quantum propositions defined by the identity and the EPR state is maximal: adding any proposition involves a contradiction.

by replacing certain classical dynamical variables, like position and momentum, with mathematical representatives—matrices—which do not commute. Shortly afterwards, Schrödinger developed a wave mechanical version of quantum mechanics and proved the formal equivalence of the two theories. It is common to understand the significance of the transition from classical to quantum mechanics in terms of 'wave-particle duality,' the idea that a quantum system like an electron, unlike a classical system like a stone, manifests itself as a wave under certain circumstances and as a particle under other circumstances. This picture obscures far more than it illuminates. We can see more clearly what is going on conceptually if we consider the implications of Heisenberg's move for the way we think about objects and their properties in the most general sense.

Heisenberg replaced the commutative algebra of dynamical variables of classical mechanics—position, momentum, angular momentum, energy, etc.—with a noncommutative algebra. Some of these dynamical variables take the values $0$ and $1$ only and correspond to properties. For example, we can represent the property of a particle being in a certain region of space by a dynamical variable that takes the value $1$ when the particle is in the region and $0$ otherwise. A dynamical variable like position corresponds to a set of such 2-valued dynamical variables or physical properties. In the case of the position of a particle, these are the properties associated with the particle being in region $R$, for all regions $R$. If, for all regions $R$, you know whether or not the particle is in that region, you know the position of the particle, and conversely. The 2-valued dynamical variables or properties of a classical system form a Boolean algebra, a subalgebra of the commutative algebra of dynamical variables.

Replacing the commutative algebra of dynamical variables with a noncommutative algebra is equivalent to replacing the Boolean algebra of 2-valued dynamical variables or properties with a non-Boolean algebra. The really essential thing about the classical mode of representation of physical systems in relation to quantum mechanics is that the properties of classical systems are represented as having the structure of a Boolean algebra or Boolean lattice. Every Boolean lattice is isomorphic to a lattice of subsets of a set.[23] To say that the properties of a classical system form a Boolean lattice is to say that they can be represented as the subsets of a set, the phase space or state space of classical mechanics. To say that a physical system has a certain property is to associate the system with a certain set in a representation space where the elements of the space—the points of the set—represent all possible states of the system. A state picks out a collection of sets, the sets to which the point representing the state belongs, as the properties of the system in that state. The dynamics of classical mechanics is described in terms of a law of motion describing how the state moves in the state space. As the state changes with time, the set of properties selected by the state changes. (For

[23] A lattice is a partially ordered set in which every pair of elements has a greatest lower bound (or infimum) and least upper bound (or supremum) with respect to the ordering, a minimum element (denoted by 0), and a maximum element (denoted by 1). A Boolean lattice is a complemented, distributive lattice, i.e., every element has a complement (the lattice analogue of set-theoretic complementation) and the distributive law holds for the infimum and supremum. The partial ordering in a Boolean lattice represented by the subsets of a set $X$ corresponds to the partial ordering defined by set inclusion, so the infimum corresponds to set intersection, the supremum corresponds to set union, 0 corresponds to the null set, and 1 corresponds to the set $x$. A Boolean algebra, defined in terms of algebraic sum ($+$) and product (.) operations, is equivalent to a Boolean lattice defined as a partially ordered structure.

an elaboration, see [Hughes, 1995] and [Bub, 1997].)

So the transition from classical to quantum mechanics involves replacing the representation of properties as a Boolean lattice, i.e., as the subsets of a set, with the representation of properties as a certain sort of non-Boolean lattice. Dirac and von Neumann developed Schrödinger's equivalence proof into a representation theory for the properties of quantum systems as subspaces in a linear vector space over the complex numbers: Hilbert space. The non-Boolean lattice in question is the lattice of subspaces of this space. Instead of representing properties as the subsets of a set, quantum mechanics represents properties as the subspaces of a linear space—as lines, or planes, or hyperplanes, i.e., as a projective geometry. Algebraically, this is the central structural change in the transition from classical to quantum mechanics—although there is more to it: notably the fact that the state space for quantum systems is a Hilbert space over the complex numbers, not the reals, which is reflected in physical phenomena associated with the possibility of superposing states with different relative phases.

Instead of talking about properties, we can talk equivalently about propositions. (We say that a given property is instantiated if and only if the corresponding proposition is true.) In a Boolean propositional structure, there exist 2-valued homomorphisms on the structure that correspond to truth-value assignments to the propositions. In fact, each point in phase space—representing a classical state—defines a truth-value assignment to the subsets representing the propositions: each subset to which the point belongs represents a true proposition or a property that is instantiated by the system, and each subset to which the point does not belong represents a false proposition or a property that is not instantiated by the system. So a classical state corresponds to a complete assignment of truth values to the propositions, or a maximal consistent 'list' of properties of the system, and all possible states correspond to all possible maximal consistent lists.

Probabilities can be introduced on such a classical property structure as measures on the subsets representing the properties. Since each phase space point defines a truth-value assignment, the probability of a property is the measure of the set of truth-value assignments that assign a 1 ('true') to the property—in effect, we 'count' (in the measure-theoretic sense) the relative number of state descriptions in which the property is instantiated (or the corresponding proposition is true), and this number represents the probability of the property. So it makes sense to interpret the probability of a property as a measure of our ignorance as to whether or not the property is instantiated. Probability distributions over classical states represented as phase space points are sometimes referred to as 'mixed states,' in which case states corresponding to phase space points are distinguished as 'pure states.'

The problem for a quantum property structure, represented by the lattice of subspaces of a Hilbert space, arises because 2-valued homomorphisms do not exist on these structures (except in the special case of a 2-dimensional Hilbert space). If we take the subspace structure of Hilbert space seriously as the structural feature of quantum mechanics corresponding to the Boolean property structure or propositional structure of classical mechanics, the non-existence of 2-valued homomorphisms on the lattice of subspaces of a Hilbert space means that there is no partition of the totality of properties of the assocated quantum system into two sets: the properties that are instantiated by the system, and the properties that are not instantiated by the system; i.e., there

is no partition of the totality of propositions into true propositions and false propositions. (Of course, other ways of associating propositions with features of a Hilbert space are possible, and other ways of assigning truth values, including multi-valued truth value assignments and contextual truth value assignments. Ultimately, the issue here concerns what we take as the salient structural change involved in the transition from classical to quantum mechanics.)

It might appear that, on the standard interpretation, a pure quantum state represented by a 1-dimensional subspace in Hilbert space—a minimal element in the subspace structure—defines a truth-value assignment on quantum propositions in an analogous sense to the truth-value assignment on classical propositions defined by a pure classical state. Specifically, on the standard interpretation, a pure quantum state selects the propositions represented subspaces containing the state as true, and the propositions represented by subspaces orthogonal to the state as false. (Note that orthogonality is the analogue of set-complement, or negation, in the subspace structure; the set-theoretical complement of a subspace is not in general a subspace.)

There is, however, an important difference between the two situations. In the case of a classical state, every possible property represented by a phase space subset is selected as either instantiated by the system or not; equivalently, every proposition is either true or false. In the case of a quantum state, the properties represented by Hilbert space subspaces are not partitioned into two such mutually exclusive and collectively exhaustive sets: some propositions are assigned no truth value. Only propositions represented by subspaces that contain the state are assigned the value 'true,' and only propositions represented by subspaces orthogonal to the state are assigned the value 'false.' This means that propositions represented by subspaces that are at some non-zero or non-orthogonal angle to the ray representing the quantum state are not assigned any truth value in the state, and the corresponding properties must be regarded as *indeterminate* or *indefinite*: according to the theory, there *can be no fact of the matter* about whether these properties are instantiated or not.

It turns out that there is only one way to assign (generalized) probabilities to quantum properties, i.e., weights that satisfy the usual Kolmogorov axioms for a probability measure on Boolean sublattices of the non-Boolean lattice of quantum properties. This is the content of Gleason's theorem [Gleason, 1957]. For a quantum state $\rho$, a property $p$ represented by a projection operator $P$ is assigned the probability $\mathrm{Tr}(\rho P)$. If $\rho$ is a pure state $\rho = |\psi\rangle\langle\psi|$, the probability of $p$ is $|\langle\psi_p|\psi\rangle|^2$, where $|\psi_p\rangle$ is the orthogonal projection of $|\psi\rangle$ onto the subspace $P$, i.e., the probability of $p$ is the square of the cosine of the angle between the ray $|\psi\rangle$ and the subspace $P$. This means that properties represented by subspaces containing the state are assigned probability 1, properties represented by subspaces orthogonal to the state are assigned probability 0, and all other properties, represented by subspaces at a non-zero or non-orthogonal angle to the state are assigned a probability between 0 and 1. So quantum probabilities are not represented as measures over truth-value assignments and cannot be given an ignorance interpretation in the obvious way.

The question now is: *what do these 'angle probabilities' or, perhaps better, 'angle weights' mean?* The orthodox answer is that the probability assigned to a property of a system by a quantum state is to be understood as the probability of finding the property in a measurement process designed to ascertain whether or not that property obtains.

A little thought will reveal that this proposal is very problematic. When the system is represented by a quantum state that assigns a certain property the probability $1/2$, say, this property is indeterminate. Physicists would say that ascribing the property to the system in that state is 'meaningless.' But somehow it makes sense to design an experiment to ascertain whether or not the property is instantiated by the system. And in such a measurement, the probability is asserted to be $1/2$ that the experiment will yield the answer 'yes,' and $1/2$ that the experiment will yield the answer 'no.' Clearly, a measurement process in quantum mechanics is not simply a procedure for ascertaining whether or not a property is instantiated in any straightforward sense. Somehow, a measurement process enables an indeterminate property, that is neither instantiated nor not instantiated by a system in a given quantum state, to either instantiate itself or not with a certain probability; or equivalently, a proposition that is neither true nor false can become true or false with a certain probability in a suitable measurement process.

The probability problem (as opposed to the truth problem, Schrödinger's problem) is the problem of interpreting the 'angle weights' as probabilities in some sense (relative frequencies? propensities? subjective Bayesian betting probabilities?) that does not reduce to a purely instrumentalist interpretation of quantum mechanics, according to which the theory is simply regarded as a remarkably accurate instrument for prediction. (Recall Einstein's remark about quantum mechanics being of interest only to shopkeepers and engineers on the Copenhagen interpretation.) The problem arises because of the unique way in which probabilities can be introduced in quantum mechanics, and because the notion of measurement or observation is utterly mysterious on the Copenhagen interpretation.

In classical theories, we measure to find out what *we don't know*, but in principle a measurement does not change *what is* (and even if it does change what is, this is simply a change or disturbance from one state of being to another that can be derived on the basis of the classical theory itself). In quantum mechanics, measurements apparently bring into being something that was *indeterminate*, not merely unknown, before, i.e., a proposition that was neither true nor false becomes true in a measurement process, and the way in which this happens according to the theory is puzzling, given our deepest assumptions about objectivity, change, and intervention.

Now, *we know how to solve Schrödinger's problem*, i.e., we know all the possible ways of modifying quantum mechanics to solve this problem. The problem arises because of the linear dynamics of the theory, which yields a certain entangled state as the outcome of a measurement interaction, and the interpretation of this entangled state as representing a state of affairs that makes certain propositions true, certain propositions false, and other propositions indeterminate. Either we change the linear dynamics in some way, or we keep the linear dynamics and say something non-orthodox about the relation between truth and indeterminateness and the quantum state. Both options have been explored in various ways and in great detail: *we understand the solution space for Schrödinger's problem, and the consequences of adopting a particular solution*.

'Collapse' theories, like the theory developed by Ghirardi, Rimini, and Weber (GRW), and extended by Pearle [Ghirardi, 2002], solve the problem by modifying the linear dynamics of quantum mechanics. (See Dickson, this vol., ch. 4, for an account.) In the modified theory, there is a certain very small probability that the wavefunction of a particle (the function defined by the components of the quantum state with respect to

the position basis in Hilbert space) will spontaneously 'collapse' after being multiplied by a peaked Gaussian of a specified width. For a macroscopic system consisting of many particles, this probability can be close to 1 for very short time intervals. In effect, this collapse solution modifies the linear dynamics of standard quantum mechanics by adding uncontrollable noise. When the stochastic terms of the modified dynamics become important at the mesoscopic and macroscopic levels, they tend to localize the wave function in space. So measurement interactions involving macroscopic pieces of equipment (or cats) can be distinguished from elementary quantum processes, insofar as they lead to the almost instantaneous collapse of the wave function and the correlation of the measured observable with the position of a localized macroscopic pointer observable.

'No collapse' solutions are constrained by certain 'no go' theorems that restrict the assignment of properties, or values to observables, under very general assumptions about the algebra of observables [Kochen and Specker, 1967], or restrict the assignment of values to observables under certain assumptions about how distributions of values are related to quantum probabilities [Bell, 1964]. A theorem by Bub and Clifton [1996] shows that if you assume that the set of definite-valued observables has a certain structure (essentially allowing quantum probabilities to be recovered as classical measures over distributions defined by different possible sets of values or properties), and the pointer observable in a measurement process belongs to the set of definite-valued observables, then the class of such theories—so-called 'modal interpretations'—is uniquely specified. More precisely, the sublattice associated with any single observable $R$ is a Boolean lattice, $\mathcal{B}$, and a quantum state $|\psi\rangle$ defines a classical probability measure on $\mathcal{B}$, in the sense that all the single and joint probabilities assigned by $|\psi\rangle$ to elements in $\mathcal{B}$ can be recovered as measures on a Kolmogorov probability space defined on the 'phase space' $X$ of 2-valued homomorphisms on $\mathcal{B}$. The Bub-Clifton theorem characterizes the maximal lattice extension, $\mathcal{L}$, of any such Boolean sublattice associated with an observable $R$ and a given quantum state $|\psi\rangle$, under the assumption that $\mathcal{L}$ is an ortholattice,[24] invariant under lattice automorphisms that preserve $R$ and $|\psi\rangle$, for which the probabilities assigned by $|\psi\rangle$ to elements in $\mathcal{L}$ can be similarly recovered as measures on a Kolmogorov probability space defined on the 'phase space' $Y$ of 2-valued homomorphisms on $\mathcal{L}$. In this sense, the theorem characterizes the limits of classicality in a quantum propositional structure. It turns out that different modal interpretations can be associated with different 'determinate sublattices' $\mathcal{L}$, i.e., with different choices of a 'preferred observable' $R$. For standard quantum mechanics, $R$ is the identity, and the determinate sublattice $\mathcal{L}$ consists of all quantum propositions represented by subspaces containing the state $|\psi\rangle$ (propositions assigned probability 1 by $|\psi\rangle$) and subspaces orthogonal to $|\psi\rangle$ (propositions assigned probability 0 by $|\psi\rangle$. Bohm's hidden variable theory can be regarded as a modal interpretation in which the preferred observable is position in configuration space. (See Dickson, this vol., ch. 4, and [Goldstein, 2001] for an account of Bohm's theory.)

An alternative type of 'no-collapse' solution to the Schrödinger problem is provided by the Everett interpretation [Everett, 1957]. (See Dickson, this vol., ch. 4, for an account.) There are a variety of Everettian interpretations in the literature, the com-

[24] I.e., an orthogonal complement exists for every element of $\mathcal{L}$.

mon theme being that all possible outcomes of a measurement are regarded as actual in some indexical sense, relative to different terms in the global entangled state (with respect to a certain preferred basis in Hilbert space), which are understood to be associated with different worlds or different minds, depending on the version. The most sophisticated formulation of Everett's interpretation is probably the Saunders-Wallace version [Saunders, 1998, Wallace, 2003]. Here the preferred basis is selected by decoherence (see below), and probabilities are introduced as rational betting probabilities in the Bayesian sense via a decision-theoretic argument originally due to Deutsch [1999].

To sum up: any solution to Schrödinger's measurement problem involves either modifying the linear dynamics of the theory ('collapse' theories), or taking some observable in addition to the identity as having a determinate value in every quantum state, and modifying what the standard theory says about what propositions are true, false, and indeterminate in a quantum state (modal interpretations, 'no collapse' hidden variable theories), so that at the end of a measurement interaction that correlates macroscopic pointer positions with possible values of a measured observable, the pointer propositions and propositions referring to measured values end up having determinate truth values. Alternatively (Everettian interpretations), we can interpret quantum mechanics so that every measurement outcome becomes determinate in some indexical sense (with respect to different worlds, or different minds, or different branches of the entangled state, etc.).

We know in considerable detail what these solutions look like, in terms of how quantum mechanics is modified. It was a useful project to explore these solutions, because we learnt something about quantum mechanics in the process, and perhaps there is more to learn by exploring the solution space further. But the point to note here is that all these solutions to the 'truth problem' of measurement distort quantum mechanics in various ways by introducing additional structural features that obscure rather than illuminate our understanding of the phenomena involved in information-theoretic applications of entanglement, such as quantum teleportation, the possibility and impossibility of certain quantum cryptographic protocols relative to classical protocols, the exponential speed-up of quantum computation algorithms relative to classical algorithms, and so on.

Consider again the Bohr-Einstein dispute about the interpretation of quantum mechanics. One might say that what separated Einstein (and Schrödinger) and Bohr was their very different answers to what van Fraassen [1991, p. 4] has called 'the foundational question *par excellence*: *how could the world possibly be the way quantum theory says it is?* This would be misleading. Einstein answered this question by arguing that the world couldn't be the way quantum theory says it is, unless the theory is not the whole story (so a 'completion' of the theory—perhaps Einstein's sought-after unified field theory—would presumably answer the question). But Bohr's complementarity interpretation is not intended to be an answer to this question. Rather, complementarity should be understood as suggesting an answer to a different question: *why must the world be the way quantum theory says it is?*

To bring out the difference between these two questions, consider Einstein's distinction between what he called 'principle' versus 'constructive' theories. Einstein introduced this distinction in an article on the significance of the special and general theories of relativity that he wrote for the London *Times*, which appeared in the issue

of November 28, 1919 [1919]:

> We can distinguish various kinds of theories in physics. Most of them are constructive. They attempt to build up a picture of the more complex phenomena out of the material of a relatively simple formal scheme from which they start out. Thus the kinetic theory of gases seeks to reduce mechanical, thermal, and diffusional processes to movements of molecules—i.e., to build them up out of the hypothesis of molecular motion. When we say that we have succeeded in understanding a group of natural processes, we invariably mean that a constructive theory has been found which covers the processes in question.
>
> Along with this most important class of theories there exists a second, which I will call 'principle theories.' These employ the analytic, not the synthetic, method. The elements which form their basis and starting-point are not hypothetically constructed but empirically discovered ones, general characteristics of natural processes, principles that give rise to mathematically formulated criteria which the separate processes or the theoretical representations of them have to satisfy. Thus the science of thermodynamics seeks by analytical means to deduce necessary conditions, which separate events have to satisfy, from the universally experienced fact that perpetual motion is impossible.

Einstein's point was that relativity theory is to be understood as a principle theory. He returns to this theme in his 'Autobiographical Notes' [1949, pp. 51–52], where he remarks that he first tried to find a constructive theory that would account for the known properties of mater and radiation, but eventually became convinced that the solution to the problem was to be found in a principle theory that reconciled the constancy of the velocity of light in vacuo for all inertial frames of reference, and the equivalence of inertial frames for all physical laws (mechanical as well as electromagnetic):

> Reflections of this type made it clear to me as long ago as shortly after 1900, i.e., shortly after Planck's trailblazing work, that neither mechanics nor electrodynamics could (except in limiting cases) claim exact validity. By and by I despaired of the possibility of discovering the true laws by means of constructive efforts based on known facts. The longer and the more despairingly I tried, the more I came to the conviction that only the discovery of a universal formal principle could lead us to assured results. The example I saw before me was thermodynamics. The general principle was there given in the theorem: the laws of nature are such that it is impossible to construct a *perpetuum mobile* (of the first and second kind). How, then, could such a universal principle be found?

A little later [1949, p. 57], he adds:

> The universal principle of the special theory of relativity is contained in the postulate: The laws of physics are invariant with respect to the Lorentz-transformations (for the transition from one inertial system to any other

> arbitrarily chosen system of inertia). This is a restricting principle for natural laws, comparable to the restricting principle for the non-existence of the *perpetuum mobile* which underlies thermodyamics.

According to Einstein, two very different sorts of theories should be distinguished in physics. One sort involves the reduction of a domain of relatively complex phenomena to the properties of simpler elements, as in the kinetic theory, which reduces the mechanical and thermal behavior of gases to the motion of molecules, the elementary building blocks of the constructive theory. The other sort of theory is formulated in terms of 'no go' principles that impose constraints on physical processes or events, as in thermodynamics ('no perpetual motion machines'). For an illuminating account of the role played by this distinction in Einstein's work, see the discussion by Martin Klein in [1967].

The special theory of relativity is a principle theory, formulated in terms of two principles: the equivalence of inertial frames for all physical laws (the laws of electromagnetic phenomena as well as the laws of mechanics), and the constancy of the velocity of light in vacuo for all inertial frames. These principles are irreconcilable in the geometry of Newtonian space-time, where inertial frames are related by Galilean transformations. The required revision yields Minkowski geometry, where inertial frames are related by Lorentz transformations. Einstein characterizes the special principle of relativity, that the laws of physics are invariant with respect to Lorentz transformations from one inertial system to another, as 'a restricting principle for natural laws, comparable to the restricting principle for the non-existence of the *perpetuum mobile* which underlies thermodynamics.' (In the case of the general theory of relativity, the group of allowable transformations includes all differentiable transformations of the space-time manifold onto itself.) By contrast, the Lorentz theory [Lorentz, 1909], which derives the Lorentz transformation from the electromagnetic properties of the aether, and assumptions about the transmission of molecular forces through the aether, is a constructive theory.

The question:

> How could the world possibly be the way the quantum theory says it is?

is motivated by a difficulty in interpreting quantum mechanics as a constructive theory, and the appropriate response is some constructive repair to the theory that resolves the difficulty, or the demonstration that the puzzling features of quantum mechanics at the phenomenal level (the phenomena of interference and entanglement) can be derived from a physically unproblematic constructive theory.

The question:

> Why must the world be the way the quantum theory says it is?

does not ask for a 'bottom-up' explanation of quantum phenomena in terms of a physical ontology and dynamical laws. Rather, the question concerns a 'top-down' derivation of quantum mechanics as a principle theory, in terms of operational constraints on the possibilities of manipulating phenomena. In the case of quantum mechanics, the relevant phenomena concern information.

This shift in perspective between the two questions is highlighted in a remark by Andrew Steane in his review article on ‘Quantum Computing’ [1998, p.119]:

> Historically, much of fundamental physics has been concerned with discovering the fundamental particles of nature and the equations which describe their motions and interactions. It now appears that a different programme may be equally important: to discover the ways that nature allows, and prevents, *information* to be expressed and manipulated, rather than particles to move.

Steane concludes his review with the following proposal [1998, p. 171]:

> To conclude with, I would like to propose a more wide-ranging theoretical task: to arrive at a set of principles like energy and momentum conservation, but which apply to information, and from which much of quantum mechanics could be derived. Two tests of such ideas would be whether the EPR-Bell correlations thus became transparent, and whether they rendered obvious the proper use of terms such as ‘measurement’ and ‘knowledge.’

A similar shift in perspective is implicit in Wheeler’s question ‘Why the quantum?,’ one of Wheeler’s ‘Really Big Questions’ [1998]. Steane’s suggestion is to answer the question by showing how quantum mechanics can be derived from information-theoretic principles. A more specific proposal along these lines originates with Gilles Brassard and Chris Fuchs. As remarked in §5.2.1, Brassard and Fuchs [Brassard, 2000, Fuchs, 1997, Fuchs, 2000, Fuchs and Jacobs, 2002] speculated that quantum mechanics could be derived from information-theoretic constraints formulated in terms of certain primitive cryptographic protocols: specifically, the possibility of unconditionally secure key distribution, and the impossibility of unconditionally secure bit commitment.

The CBH theorem (motivated by the Brassard-Fuchs conjecture) shows that quantum mechanics can be regarded as a principle theory in Einstein’s sense, where the principles are information-theoretic constraints. So we have an answer to the question: why must the world be the way quantum mechanics says it is? The phenomena of interference and nonlocal entanglement are bound to occur in a world in which there are certain constraints on the acquisition, communication, and processing of information.

Consider, for comparison, relativity theory, the other pillar of modern physics. A relativistic theory is a theory with certain symmetry or invariance properties, defined in terms of a group of space-time transformations. Following Einstein’s formulation of special relativity as a principle theory, we understand this invariance to be a consequence of the fact that we live in a world in which natural processes are subject to certain constraints: roughly (as Hermann Bondi [1980] puts it), ‘no overtaking of light by light,’ and ‘velocity doesn’t matter’ (for electromagnetic as well as mechanical phenomena). Recall Einstein’s characterization of the special principle of relativity as ‘a restricting principle for natural laws, comparable to the restricting principle of the non-existence of the *perpetuum mobile* which underlies thermodynamics.’) Without Einstein’s analysis, the transformations of Minkowski space-time would simply be a rather puzzling algorithm for relativistic kinematics and the Lorentz transformation, which is

incompatible with the kinematics of Newtonian space-time. What Einstein's analysis provides is a rationale for taking the structure of space and time as Minkowskian: we see that this is required for the consistency of the two principles of special relativity.

A quantum theory is a theory in which the observables and states have a certain characteristic algebraic structure. Unlike relativity theory, quantum mechanics was born as a recipe or algorithm for caclulating the expectation values of observables measured by macroscopic measuring instruments. A theory with a commutative $C^*$-algebra has a phase space representation—not necessarily the phase space of classical mechanics, but a theory in which the observables of the $C^*$-algebra are replaced by 'beables' (Bell's term, see [1987]), and the $C^*$-algebraic states are replaced by beable-states representing complete lists of properties (idempotent quantities). In this case, it is possible to extend the theory to include the measuring instruments that are the source of the $C^*$-algebraic statistics, so that they are no longer 'black boxes' but constructed out of systems that are characterized by properties and states of the phase space theory. That is, the $C^*$-algebraic theory can be replaced by a 'detached observer' theory of the physical processes underlying the phenomena, to use Pauli's term [Born, 1971, p. 218], including the processes involved in the functioning of measuring instruments. Note that this depends on a representation theorem. In the noncommutative case, we are guaranteed only the existence of a Hilbert space representation of the $C^*$-algebra, and it is an open question whether a 'detached observer' description of the phenomena is possible.

Solving Schrödinger's problem—the truth problem—amounts to a proposal to treat quantum mechanics as a failed or incomplete constructive theory in need of constructive repair. In effect, the problem is *how to account for quantum information*—the puzzling features of interference and nonlocal entanglement—in a theoretical framework in which only classical information is meaningful in a fundamental sense. If we treat quantum mechanics as a principle theory of information, the core foundational problem is the probability problem. From this perspective, the problem is *how to account for the appearance of classical information* in a quantum world characterized by information-theoretic constraints.

One might complain that treating quantum mechanics as a principle theory amounts to simply postulating what is ultimately *explained* by a constructive theory like the GRW theory or Bohm's theory. This would amount to rejecting the idea that a principle theory can be explanatory. From the perspective adopted here, Bohm's constructive theory in relation to quantum mechanics is like Lorentz's constructive theory of the electron in relation to special relativity. Cushing [1998, p. 204] quotes Lorentz (from the conclusion of the 1916 edition of *The Theory of Electrons*) as complaining similarly that 'Einstein simply postulates what we have deduced.'

> I cannot speak here of the many highly interesting applications which Einstein has made of this principle [of relativity]. His results concerning electromagnetic and optical phenomena ... agree in the main with those which we have obtained in the preceding pages, the chief difference being that Einstein simply postulates what we have deduced, with some difficulty and not altogether satisfactorily, from the fundamental equations of the electromagnetic field. By doing so, he may certainly take credit for

> making us see in the negative result of experiments like those of Michelson, Rayleigh and Brace, not a fortuitous compensation of opposing effects, but the manifestation of a general and fundamental principle.
>
> Yet, I think, something may also be claimed in favour of the form in which I have presented the theory. I cannot but regard the aether, which can be the seat of an electromagnetic field with its energy and its vibrations, as endowed with a certain degree of substantiality, however different it may be from all ordinary matter. In this line of thought, it seems natural not to assume at starting that it can never make any difference whether a body moves through the aether or not, and to measure distances and lengths of time by means of rods and clocks having a fixed position relative to the aether.

Note that Lorentz's theory is constrained by the principles of special relativity, which means that the aether as a rest frame for electromagnetic phenomena must, in principle, be undetectable. So such a theory can have no excess empirical content over special relativity. Cushing [1998, p. 193] also quotes Maxwell as asking whether 'it is not more philosophical to admit the existence of a medium which we cannot at present perceive, than to assert that a body can act at a place where it is not.' Yes, but not if we also have to admit that, in principle, as a matter of physical law, if we live in a world in which events are constrained by the two relativistic principles, the medium must remain undetectable.

You can, if you like, tell a constructive story about quantum phenomena, but such an account, if constrained by the information-theoretic principles, will have no excess empirical content over quantum mechanics. Putting this differently, a solution to Schrödinger's truth problem that has excess empirical content over quantum mechanics must violate one or more of the CBH information-theoretic constraints. So, e.g., a Bohmian theory of quantum phenomena is like an aether theory for electromagnetic fields. Just as the aether theory attempts to make sense of the behaviour of fields by proposing an aether that is a sort of *sui generis* mechanical system different from all other mechanical systems, so Bohm's theory attempts to make sense of quantum phenomena by introducing a field (the quantum potential or guiding field) that is a sort of *sui generis* field different from other physical fields.

The crucial distinction here is between a constructive theory formulated in terms of a physical ontology and dynamical laws ('bottom-up') and a principle theory formulated in terms of operational constraints at the phenomenal level ('top-down'). A constructive theory introduces an algebra of beables and beable-states. A principle theory introduces an algebra of observables and observable-states, which are essentially probability measures.

It seems clear that the algebra of observables will be non-trivially distinct from the algebra of beables if cloning is impossible. For if a constructive theory for a certain domain of phenomena allows dynamical interactions in which a beable of one system, designated as the measuring instrument, can become correlated with a beable of another system, designated as the measured system, without disturbing the values of other beables of the measured system, we can take such an interaction as identifying the value of the beable in question (in the sense that the value of a beable of one system

is recorded in the value of a beable of a second system). If this is possible, then it will be possible to simultaneously measure any number of beables of a system by concatenating measurement interactions, and so it will be possible in principle to identify any arbitrary beable state. If we assume that we can prepare any state, then the possibility of identifying an arbitrary state means that we can construct a device that could copy any arbitrary state. So if we *cannot* construct such a device, then measurement in this sense must also be impossible. It follows that a 'measurement' in the constructive theory will be something other than the mere identification of a beable value of a system, without disturbance, and the question of what the observables are in such a theory will require a non-trivial analysis.

Such an analysis is indeed given by Bohm in Part II of his two-part 1952 paper on hidden variables [1952], and a more careful and sophisticated analysis is given by [Dürr *et al.*, 2003] for their 'Bohmian mechanics' version of Bohm's theory. As one would expect (given the equilibrium distribution assumption which ensures that Bohm's theory is empirically indistinguishable from quantum mechanics), while the *beables* are functions of position in configuration space (and form a commutative algebra), the *observables* of the theory are just the observables of quantum theory and form a noncommutative algebra.

The CBH theorem assumes that, for the theories we are concerned with, the *observables* form a $C^*$-algebra. The content of the CBH theorem is that, given certain information-theoretic constraints, the $C^*$-algebra of *observables* and *observable states* takes a certain form characteristic of quantum theories. The theorem says nothing about beables and beable-states, and does not address the measurement problem (Schrödinger's truth problem), let alone solve it. But from the perspective adopted here, the measurement problem is simply the observation that cloning is impossible, and a 'solution to the measurement problem' is the proposal of a physical ontology and dynamics and an analysis of measurement that yields the observables and observable-states of standard quantum mechanics. Such theories provide possible explanations for the impossibility of cloning. But since there are now a variety of such explanations available, and—assuming the CBH information-theoretic principles—there are no empirical constraints, in principle, that could distinguish these explanations, there seems little point in pursuing the question further. A constructive theory whose *sole* motivation is to 'solve the measurement problem' seems unlikely to survive fundamental advances in physics driven by other theoretical or experimental problems

The probability problem—the core foundational problem for the interpretation of quantum mechanics as a principle theory of information—can be put this way: From the information-theoretic constraints, we get a noncommutative (or non-Boolean) theory of correlations for which there is no phase space representation. One can define, in a unique way (according to Gleason's theorem) generalized 'transition probabilities' or 'transition weights' associated with certain structural features of the noncommutative structure: the angles between geometrical elements representing quantum 'propositions.' The problem is how to understand these weights as representing probabilities, *without reducing the problem to a solution of the truth problem.*

It seems clear that we need to take account of the phenomenon of decoherence (see Landsmann, this vol., ch. 5; Dickson, this vol., ch. 4; [Zurek, 2003, Olliver *et al.*, 2004]): an extremely fast process that occurs in the spontaneous interaction between a macrosys-

tem and its environment that leads to the virtually instantaneous suppression of quantum interference. What happens, roughly, is that a macrosystem like Schrödinger's cat typically becomes correlated with the environment—an enormous number of stray dust particles, air molecules, photons, background radiation, etc.—in an entangled state that takes a certain form with respect to a preferred set of basis states, which remain stable as the interaction develops and includes more and more particles. It is as if the environment is 'monitoring' the system via a measurement of properties associated with the preferred states, in such a way that information about these properties is stored redundantly in the environment. This stability, or robustness, of the preferred basis, and the redundancy of the information in the environment, allows one to identify certain emergent structures in the overall pattern of correlations—such as macroscopic pointers and cats and information-gatherers in general—as classical-like: the correlational information required to reveal quantum interference for these structures is effectively lost in the environment. So it appears that the information-theoretic constraints are consistent with both (i) the conditions for the existence of measuring instruments as sources of classical information, and (ii) the existence of information-gatherers with the ability to use measuring instruments to apply and test quantum mechanics, given a characterization of part of the overall system as the environment. That is, decoherence provides an explanation for the emergence of classical information in a quantum correlational structure.

If something like the above account of decoherence is acceptable, then the probability problem reduces to showing that the probabilities assigned to measurement outcomes by these information-gatherers, in the subjective Bayesian sense, are just the Gleason generalized transition probabilities. That is, we need to show that, while quantum theory, at the fundamental level, is a noncommutative theory of correlations for which there is no phase space representation, it is *also* a theory of the probabilistic behavior of information-gatherers, certain emergent structures in the pattern of correlations when correlational information in their environment is ignored. For an argument along these lines, see [Pitowsky, 2002].

On the view proposed here, no measurement outcomes are certified as determinate by the theory. Rather, measuring instruments are sources of classical information, where the individual occurrence of a particular distinguishable event ('symbol') produced stochastically by the information source lies outside the theory. In this sense, a measuring instrument, insofar as it functions as a classical information source, is still ultimately a 'black box' in the theory. So a quantum description will have to introduce a 'cut' between what we take to be the ultimate measuring instrument in a given measurement process and the quantum phenomenon revealed by the instrument. But this 'cut' is no longer *ad hoc*, or mysterious, or in some other way problematic, as it is in the Copenhagen interpretation (see Landsmann, this vol., ch. 5). For here the 'cut' just reflects the fundamental interpretative claim: that quantum mechanics is *a theory about the representation and manipulation of information* constrained by the possibilities and impossibilities of information-transfer in our world, rather than a theory about the ways in which nonclassical waves and particles move.

# References

[Aharonov *et al.*, 1964] Yakir Aharonov, Paul Gabriel Bergmann, and Jerome L. Lebowitz. Time-symmetric quantum states. *Physical Review B*, 134:1410–1416, 1964. Reprinted in *Quantum Theory and Measurement*, John Archibald Wheeler and Wojciech H. Zurek (eds.), Princeton University Press, Princeton (1983), pp. 680–686.

[Aharonov *et al.*, 2005] Dorit Aharonov, Amnon Ta-Shma, Umesh V. Vazirani, and Andrew C. Yao. Quantum bit escrow. In *STOC 2000, Proceedings of the Thirty Second Annual ACM Symposium on Theory Of Computing. arXiv e-print quant-ph/0004017*, 2005.

[Aspect *et al.*, 1981] Alain Aspect, Phillipe Grangier, and G. Roger. Experimental tests of realistic local theories via Bell's theorem. *Physical Review Letters*, 47:460–477, 1981.

[Aspect *et al.*, 1982] Alain Aspect, Phillipe Grangier, and G. Roger. Experimental realization of EPR gedankenexperiment: a new violation of Bell's inequalities. *Physical Review Letters*, 49:91–94, 1982.

[Bacciagaluppi, 1994] Guido Bacciagaluppi. Separation theorems and Bell inequalities in algebraic quantum mechanics. In Paul Busch, Pekka Lahti, and Peter Mittelstaedt, editors, *Symposium on the Foundations of Modern Physics 1993: Quantum Measurement, Irreversibility and the Physics of Information*, pages 29–37. World Scientific, Singapore, 1994.

[Barenco, 1998] Adriano Barenco. Quantum computation: an introduction. In Lo et al. [1998], pages 143–183.

[Barnum *et al.*, 1996a] Howard Barnum, Carlton M. Caves, Christopher A. Fuchs, Richard Jozsa, and Benjamin Schumacher. Noncommuting mixed states cannot be broadcast. *Physical Review Letters*, 76:2828, 1996.

[Barnum *et al.*, 1996b] Howard Barnum, Christopher A. Fuchs, Richard Jozsa, and Benjamin Schumacher. General fidelity limit for quantum channels. *Physical Review A*, 54:4707, 1996.

[Barnum *et al.*, 2001] Howard Barnum, Patrick Hayden, Richard Jozsa, and Andreas Winter. On the reversible extraction of classical information from a quantum source. *Proceedings of the Royal Society (London) A*, 457:2019–2039, 2001.

[Bell, 1964] John Stuart Bell. On the Einstein-Podolsky-Rosen Paradox. *Physics*, 1:195–200, 1964. Reprinted in John Stuart Bell, *Speakable and Unspeakable in Quantum Mechanics*, Cambridge University Press, Cambridge, 1989.

[Bell, 1987] John Stuart Bell. Beables for quantum field theory. In *Speakable and Unspeakable in Quantum Mechanics*, pages 173–180. Cambridge University Press, Cambridge, 1987.

[Benioff, 1980] Paul Benioff. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by turing machines. *Journal of Statistical Physics*, 22(5):563–591, 1980.

[Bennett and Brassard, 1984] Charles H. Bennett and Gilles Brassard. Quantum cryptography: public key distribution and coin tossing. In *Proceedings of IEEE international conference on computers, systems, and signal processing*, pages 175–179, New York, 1984. IEEE.

[Bennett and Wiesner, 1992] Charles H. Bennett and Stephen J. Wiesner. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Physical Review Letters*, 69:2881–2884, 1992.

[Bennett *et al.*, 1982] Charles H. Bennett, Gilles Brassard, S. Breidbart, and Stephen J. Wiesner. Quantum cryptography, or unforgeable subway tokens. In D. Chaum, R.L. Rivest, and A.T. Sherman, editors, *Advances in Cryptology: Proceedings of Crypto 82*, pages 267–275, New York, 1982. Plenum Press.

[Bennett *et al.*, 1993] Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William Wootters. Teleporting an unknown quantum state via dual classical and EPR channels. *Physical Review Letters*, 70:1895–1899, 1993.

[Bennett, 1973] Charles H. Bennett. Logical reversibility of computations. *IBM Journal of Research and Development*, 17:525–532, 1973.

[Bohm, 1951] David Bohm. *Quantum Theory*. Prentice Hall, Englewood Cliffs, NJ, 1951.

[Bohm, 1952] David Bohm. A suggested interpretation of quantum theory in terms of ‘hidden’ variables. i and ii. *Physical Review*, 85:166–193, 1952.

[Bohr, 1949] Niels Bohr. Discussion with Einstein on epistemological problems in modern physics. In Paul Allen Schilpp, editor, *Albert Einstein: Philosopher-Scientist*, volume VII. The Library of Living Philosophers, Open Court, La Salle, IL, 1949.

[Bondi, 1980] Hermann Bondi. *Relativity and Common Sense*. Dover Publications, 1980.

[Born, 1971] Max Born. *The Born-Einstein Letters*. Walker and Co., London, 1971.

[Boschi *et al.*, 1998] D. Boschi, S. Branca, Francesco De Martini, Lucien Hardy, and Sandu Popescu. Experimental realization of teleporting an unknown pure quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters*, 80:1121–1125, 1998.

[Bouwmeester *et al.*, 1997] Dirk Bouwmeester, Jian-Wei Pan, K. Mattle, Manfred Eible, Harald Weinfurter, and Anton Zeilinger. Experimental quantum teleportation. *Nature*, 390:575–579, 1997.

[Brassard *et al.*, 1993] Gilles Brassard, Claude Crépeau, Richard Jozsa, and D. Langlois. A quantum bit commitment scheme provably unbreakable by both parties. In *Proceedings of the 34th Annual IEEE Symposium and the Foundations of Computer Science*, pages 362–371, New York, November 1993. IEEE.

[Brassard, 2000] Gilles Brassard. Remarks on quantum foundations in the light of quantum cryptography. 'Quantum Foundations in the Light of Quantum Information and Cryptography,' Université de Montréal, May 17–19 (2000)., 2000.

[Brukner and Zeilinger, 2000] Časlav Brukner and Anton Zeilinger. Quantum measurement and Shannon information, a reply to M.J.W. Hall. *arXiv e-print quant-ph/0008091*, 2000.

[Brukner and Zeilinger, 20002] Časlav Brukner and Anton Zeilinger. Information and fundamental elements of the structure of quantum theory. In *Festschrift for C. F. v. Weizsäcker on the occasion of his 90th birthday. arXiv e-print quant-ph/0212084*. Unknown, 20002.

[Brukner and Zeilinger, 2001] Časlav Brukner and Anton Zeilinger. Conceptual inadequacy of the Shannon information in quantum measurement. *Physical Review A*, 63:022113, 2001.

[Bub and Clifton, 1996] Jeffrey Bub and Rob Clifton. A uniqueness theorem for 'no collapse' interpretations of quantum mechanics. *Studies in the History and Philosophy of Modern Physics*, 27:181–219, 1996.

[Bub, 1997] Jeffrey Bub. *Interpreting the Quantum World*. Cambridge University Press, Cambridge, 1997.

[Bub, 2001a] Jeffrey Bub. The bit commitment theorem. *Foundations of Physics*, 3:735–756, 2001.

[Bub, 2001b] Jeffrey Bub. Secure key distribution via pre- and post-selected quantum states. *Physical Review A*, 63:032309–032311, 2001.

[Bub, 2004] Jeffrey Bub. Why the quantum? *Studies in History and Philosophy of Modern Physics*, 35B:241–266, 2004.

[Bub, 2005] Jeffrey Bub. Quantum mechanics is about quantum information. *Foundations of Physics*, 34:541–560, 2005.

[Cleve *et al.*, 1998] Richard Cleve, Arthur Ekert, Christina Macchiavello, and M. Mosca. Quantum algorithms revisited. *Proceedings of the Royal Society A*, 454:339–354, 1998.

[Clifton *et al.*, 2003] Rob Clifton, Jeffrey Bub, and Hans Halvorson. Characterizing quantum theory in terms of information-theoretic constraints. *Foundations of Physics*, 33:1561–1591, 2003.

[Connes, 1994] Alain Connes. *Noncommutative Geometry*. Academic Press, San Diego, 1994.

[Cover and Thomas, 1991] T.M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, New York, 1991.

[Cushing, 1998] James T. Cushing. *Philosophical Concepts in Physics*. Cambridge University Press, Cambridge, 1998.

[Deutsch and Jozsa, 1992] David Deutsch and Richard Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London A*, 439:553, 1992.

[Deutsch, 1985] David Deutsch. Quantum theory, the Church-Turing principle and the universal quantum compute. *Proceedings of the Royal Society of London A*, 400:497, 1985.

[Deutsch, 1989] David Deutsch. Quantum computational networks. *Proceedings of the Royal Society of London A*, 400:73, 1989.

[Deutsch, 1999] David Deutsch. Quantum theory of probability and decisions. *Proceedings of the Royal Society of London A*, 455:3129–3197, 1999.

[Dieks, 1982] Dennis Dieks. Communication by EPR devices. *Physics Letters A*, 92:271–272, 1982.

[Dürr *et al.*, 2003] Detleff Dürr, Sheldon Goldstein, and Nino Zanghì. Quantum equilibrium and the role of operators as obserables in quantum theory. *arXiv e-print quant-ph/0308038*, 2003.

[Einstein *et al.*, 1935] Albert Einstein, Boris Podolsky, and Nathan Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 47:777–780, 1935.

[Einstein, 1919] Albert Einstein. What is the theory of relativity? *The London Times*, page 13, 1919. First published November 28, 1919. Also in A. Einstein, *Ideas and Opinions*, Bonanza Books, New York, 1954, pp. 227–232.

[Einstein, 1949] Albert Einstein. Autobiographical notes. In Paul Allen Schilpp, editor, *Albert Einstein: Philosopher-Scientist*, pages 3–94. Open Court, La Salle, IL, 1949.

[Einstein, 1967] Albert Einstein. Letter to Schrödinger dated december 22, 1950. In K. Przibram, editor, *Letters on Wave Mechanics*. Philosophical Library, New York, 1967.

[Ekert, 1991] Artur Ekert. Quantum cryptography based on bell's theorem. *Physical Review Letters*, 67:661, 1991.

[Everett, 1957] Hugh Everett. 'relative state' formulation of quantum mechanics. *Reviews of Modern Physics*, 29:454–462, 1957.

[Feynman, 1982] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21:467, 1982.

[Fuchs and Jacobs, 2002] Christopher A. Fuchs and Kurt Jacobs. An information-theoretic tradeoff relation for finite-strength quantum measurements. *Physical Review A*, 63:062305–062320, 2002.

[Fuchs, 1997] Christopher A. Fuchs. Information gain vs. state disturbance in quantum theory. *Fortschr. Phys.*, 46:535–565, 1997. Reprinted in *Quantum Computation: Where Do We Want to Go Tomorrow?*, S.L. Braunstein (ed.), pp. 229–259, Wiley VCH, Weinheim (1998).

[Fuchs, 2000] Christopher A. Fuchs. Just two nonorthogonal quantum states. In *Quantum Communication, Computing, and Measurement 2*, pages 11–16, Dordrecht, 2000. Kluwer.

[Fuchs, 2001a] Christopher A. Fuchs. Notes on a Paulian idea: foundational, historical, anecdotal and forward-looking thoughts on the quantum. *arXiv e-print quant-ph/0105039*, 2001.

[Fuchs, 2001b] Christopher A. Fuchs. Quantum foundations in the light of quantum information. In A. Gonis, editor, *Proceedings of the NATO Advanced Research Workshop on Decoherence and its Implications in Quantum Computation and Information Transfer. arXiv e-print quant-ph/0106166*, 2001.

[Fuchs, 2002a] Christopher A. Fuchs. The anti-Växjo interpretation of quantum mechanics. *arXiv e-print quant-ph/0204146*, 2002.

[Fuchs, 2002b] Christopher A. Fuchs. Quantum mechanics as quantum information (and only a little more). *arXiv e-print quant-ph/0205039*, 2002.

[Furasawa *et al.*, 1998] A. Furasawa, J.L. Sorensen, Samuel L. Braunstein, Christopher A. Fuchs, H.J. Kimble, and E.S. Polzik. Unconditional quantum teleportation. *Science*, 282:706–709, 1998.

[Furry, 1936] Wendell H. Furry. A note on the quantum mechanical theory of measurement. *Physical Review*, 49:393–399, 1936.

[Ghirardi, 2002] Gian-Carlo Ghirardi. Collapse theories. In E.N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, 2002. http://plato.stanford.edu/entries/qm-collapse/.

[Gisin, 1989] Nicolas Gisin. Stochastic quantum dynamics and relativity. *Helvetica Physica Acta*, 62:363–371, 1989.

[Gleason, 1957] A.N. Gleason. Measures on the closed sub-spaces of hilbert spaces. *Journal of Mathematics and Mechanics*, 6:885–893, 1957.

[Goldstein, 2001] Sheldon Goldstein. Bohmian mechanics. In E.N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, 2001. http://plato.stanford.edu/entries/qm-bohm.

[Grover, 1997] L.K. Grover. Quantum mechanics helps in searching for a needle in a haystack. *Physical Review Letters*, 79:325, 1997.

[Hall, 2000] Michael J.W. Hall. Comment on ‘Conceptual inadequacy of the Shannon information’ by Č. Brukner and A. Zeilinger. *arXiv e-print quant-ph/0007116*, 2000.

[Halvorson and Bub, 2005] Hans Halvorson and Jeffrey Bub. Can quantum cryptography imply quantum mechanics? reply to Smolin. *Quantum Information and Computation*, 5:170–175, 2005.

[Halvorson, 2004] Hans Halvorson. Remote preparation of arbitrary ensembles and quantum bit commitment. *Journal of Mathematical Physics*, 45:4920–4931, 2004.

[Hardy and Kent, 2004] Lucien Hardy and Adrien Kent. Cheat sensitive quantum bit commitment. *Physical Review Letters*, 92:157901, 2004. Longer version available at arXiv e-print quant-ph/9911043.

[Hartley, 1928] R.V.L. Hartley. Transmission of information. *Bell System Technical Journal*, 7:53, 1928.

[Howard, 2004] Don Howard. Who invented the Copenhagen interpretation? a study in mythology. http://www.nd.edu/ dhoward1/Papers.html, 2004.

[Hughes, 1995] Richard I.G. Hughes. *The Structure and Interpretation of Quantum Mechanics*. Cambridge University Press, 1995.

[Hughston *et al.*, 1993] Lane P. Hughston, Richard Jozsa, and William K. Wootters. A complete classification of quantum ensembles having a given density matrix. *Physics Letters A*, 183:14–18, 1993.

[Jaynes, 1957] Edwin T. Jaynes. Information theory and statistical mechanics ii. *Physical Review*, 108:171–190, 1957.

[Jozsa and Linden, 2002] Richard Jozsa and Noah Linden. On the role of entanglement in quantum computational speed-up. *arXiv e-print quant-ph/0201143*, 2002.

[Jozsa and Schumacher, 1994] Richard Jozsa and Benjamin Schumacher. A new proof of the quantum noiseless coding theorem. *Journal of Modern Optics*, 41:2343–2349, 1994.

[Jozsa, 1997a] Richard Jozsa. Entanglement and quantum computation. In S.A. Huggett, L.J. Mason, K.P. Tod, S.T. Tsou, and N.M.J. Woodhouse, editors, *The Geometric Universe: Science, Geometry, and the Work of Roger Penrose*, pages 369–378. Oxford University Press, 1997.

[Jozsa, 1997b] Richard Jozsa. Quantum algorithms and the fourier transform. *arXiv e-print quant-ph/9707033*, 1997.

[Jozsa, 1998] Richard Jozsa. Quantum information and its properties. In Lo et al. [1998], pages 49–75.

[Jozsa, 1999] Richard Jozsa. Searching in grover’s algorithm. *arXiv e-print quant-ph/9901021*, 1999.

[Jozsa, 2005] Richard Jozsa. An introduction to measurement based quantum computation. *arXiv e-print quant-ph/0508124*, 2005.

[Kadison and Ringrose, 1997] Richard Kadison and J. Ringrose. *Fundamentals of the Theory of Operator Algebras*. American Mathematical Society, Providence, Rhode Island, 1997.

[Kent, 1999a] Adrian Kent. Secure classical bit commitment using fixed capacity communication channels. *arXiv e-print quant-ph/9906103*, 1999.

[Kent, 1999b] Adrian Kent. Unconditionally secure bit commitment. *Physical Review Letters*, 83:1447–1450, 1999.

[Klein, 1967] Martin J. Klein. Thermodynamics in Einstein's thought. *Science*, 157:509–516, 1967.

[Kochen and Specker, 1967] Simon Kochen and E.P. Specker. On the problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics*, 17:59–87, 1967.

[Landau, 1987] L.J. Landau. On the violation of Bell's inequality in quantum theory. *Physics Letters A*, 120:54–56, 1987.

[Landsman, 1998] Nicolaas Landsman. *Mathematical Topics Between Classical and Quantum Mechanics*. Springer, New York, 1998.

[Lewis and Papadimitriou, 1981] Harry R. Lewis and Christos H. Papadimitriou. *Elements of the theory of computation*. Prentice-Hall, Englewood Cliffs, New Jersey, 1981.

[Lo and Chau, 1997] Hoi-Kwong Lo and H.F. Chau. Is quantum bit commitment really possible? *Physical Review Letters*, 78:3410–3413, 1997.

[Lo and Chau, 1998] Hoi-Kwong Lo and H.F. Chau. Why quantum bit commitment and ideal coin tossing are impossible. *Physica D*, 120:177–187, 1998.

[Lo *et al.*, 1998] Hoi-Kwong Lo, Sandu Popescu, and Tim Spiller, editors. *Introduction to Quantum Computation and Information*, Singapore, 1998. World Scientific.

[Lorentz, 1909] H.A. Lorentz. *The Theory of Electrons*. Columbia University Press, New York, 1909.

[Mackey, 1963] George Mackey. *Mathematical Foundations of Quantum Mechanics*. Benjamin, New York, 1963.

[Mayers, 1996a] Dominic Mayers. The trouble with quantum bit commitment. *arXiv e-print quant-ph/9603015*, 1996.

[Mayers, 1996b] Dominic Mayers. Unconditionally secure quantum bit commitment is impossible. In *Proceedings of the Fourth Workshop on Physics and Computation*, pages 224–228, Boston, 1996. New England Complex System Institute.

[Mayers, 1997] Dominic Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical Review Letters*, 78:3414–3417, 1997.

[Metzger, 2000] Steffen Metzger. Spin-measurement retrodiction revisited. *arXiv e-print quant-ph/0006115*, 2000.

[Nielsen and Chuang, 2000] Michael A. Nielsen and Isaac Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, 2000.

[Nielsen *et al.*, 1998] Michael A. Nielsen, E. Knill, and Raymond Laflamme. Complete quantum teleportation using nuclear magnetic resonance. *Nature*, 396:52–55, 1998.

[Nielsen, 2003] Michael A. Nielsen. Journal club notes on cluster-state quantum computation. http://qinfo.org/qc-by-measurement/, 2003.

[Nielsen, 2005] Michael A. Nielsen. Cluster-state quantum computation. *arXiv e-print quant-ph/0504097*, 2005.

[Nyquist, 1924] H. Nyquist. Certain factors affecting telegraph speed. *Bell System Technical Journal*, 3:324, 1924.

[Olliver *et al.*, 2004] Harold Olliver, David Poulin, and W. Zurek. Environment as a witness: selective proliferation of information and emergence of objectivity in a quantum universe. *Physical Review Letters*, 93:22041, 2004.

[Pitowsky, 2002] Itamar Pitowsky. Betting on the outcomes of measurements: a Bayesian theory of quantum probability. *arXiv e-print quant-ph/0208121*, 2002.

[Popescu and Rohrlich, 1998] Sandu Popescu and David Rohrlich. The joy of entanglement. In Lo et al. [1998], pages 29–48.

[Preskill, 2005] John Preskill. Lecture notes for quantum computation. http://www.theory.caltech.edu/people/preskill/ph229/, 2005.

[Raussendorf and Briegel, 2001a] Robert Raussendorf and Hans J. Briegel. A one-way quantum computer. *arXiv e-print quant-ph/0208121*, 2001.

[Raussendorf and Briegel, 2001b] Robert Raussendorf and Hans J. Briegel. Quantum computing via measurements only. *Physical Review Letters*, 86:5188–5199, 2001.

[Rivest *et al.*, 1978] R.L. Rivest, A. Shamir, and L.M. Adleman. A method of obtaining digital signatures and public-key cryptosystems. *Communications of the Association for Computing Machinery*, 21:120–126, 1978.

[Saunders, 1998] Simon Saunders. Time, quantum mechanics, and probability. *Synthese*, 114:373–404, 1998.

[Schilpp, 1949] Paul Allen Schilpp, editor. *Albert Einstein: Philosopher-Scientist*, volume VII, La Salle, IL, 1949. The Library of Living Philosophers, Open Court.

[Schrödinger, 1935] Erwin Schrödinger. Discussion of probability relations between separated systems. *Proceedings of the Cambridge Philosophical Society*, 31:555–563, 1935.

[Schrödinger, 1936] Erwin Schrödinger. Probability relations between separated systems. *Proceedings of the Cambridge Philosophical Society*, 32:446–452, 1936.

[Schumacher, 1995] Benjamin Schumacher. Quantum coding. *Physical Review A*, 51:2738–2747, 1995.

[Schumacher, 1998] Benjamin Schumacher. Lecture notes on quantum information. University of Innsbruck, 1998.

[Shannon, 1948] Claude E. Shannon. A mathematical theory of computation. *Bell System Technical Journal*, 27:379–423, 623–656, 1948.

[Shor, 1994] Peter W. Shor. Algorithms for quantum computation; discrete logarithms and factoring. In *Proceedings, 35th Annual Symposium on Foundations of Computer Science*, Los Alamitos, CA, 1994. IEEE Press.

[Shor, 1997] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal of Computation*, 26:1484–1509, 1997.

[Simon, 1994] D.R. Simon. On the power of quantum computation. In *Proceedings, 35th Annual Symposium on Foundations of Computer Science*, pages 116–123, Los Alamitos, CA, 1994. IEEE Press.

[Simon, 1997] D.R. Simon. On the power of quantum computation. *SIAM Journal of Computation*, 26:1474–1483, 1997.

[Steane, 1998] Andrew Steane. Quantum computing. *Reports on Progress in Physics*, 61:117–173, 1998.

[Steane, 2003] Andrew Steane. A quantum computer only needs one universe. *Studies in History and Philosophy of Modern Physics*, 34B:469–478, 2003.

[Summers, 1990] Stephen J. Summers. On the independence of local algebras in quantum field theory. *Reviews in Mathematical Physics*, 2:201–247, 1990.

[Timpson, 2004] Christopher Timpson. *Quantum Information Theory and the Foundations of Quantum Mechanics*. PhD thesis, Oxford, 2004.

[Vaidman *et al.*, 1987] Lev Vaidman, Yakir Aharonov, and David Z Albert. How to ascertain the values of $s_x, s_y, s_z$ of a spin-1/2 particle. *Phys. Rev. Letters*, 58:1385, 1987.

[van Fraassen, 1991] Bas van Fraassen. *Quantum Mechanics: An Empiricist View*. Clarendon Press, Oxford, 1991.

[Vidal, 2003] Giufró Vidal. Efficient classical simulation of slightly entangled quantum computations. *Physical Review Letters*, 91:147902, 2003.

[von Neumann, 1955] John von Neumann. *Mathematical foundations of quantum mechanics*. Princeton University Press, Princeton, 1955.

[Wald, 1984] Robert M. Wald. *General Relativity*. University of Chicago Press, 1984.

[Wallace, 2003] David Wallace. Everettian rationality: defending Deutsch's approach to probability in the Everett interpretation. *Studies in the History and Philosophy of Modern Physics*, 34:415–439, 2003.

[Wheeler, 1998] John A. Wheeler. *Geons, Black Holes, and Quantum Foam: A Life in Physics*. W.W. Norton, New York, 1998. With K. Ford.

[Wiesner, 1983] Stephen J. Wiesner. Conjugate coding. *SIGACT News*, 15:17, 1983.

[Wooters and Zurek, 1982] William K. Wooters and Wojciech H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982.

[Yuen, 2005] Horace P. Yuen. Unconditionally secure quantum bit commitment. *arXiv e-print quant-ph/0505132*, 2005.

[Zurek, 2003] Wojciech H. Zurek. Decoherence, einselection, and the quantum origins of the classical. *Review of Modern Physics*, 75:715, 2003.